SAN JUAN – How It Works: Understanding DNS Abuse
Monday, March 12, 2018 – 13:30 to 15:00 AST
ICANN61 | San Juan, Puerto Rico

UNIDENTIFIED MALE: Good afternoon. ICANN61. March 12. How It Works: Understanding DNS Abuse.

CATHY PETERSEN: Good afternoon, everyone. We will be starting shortly How It Works: Understanding DNS Abuse in a couple minutes. Thank you.

Good afternoon again, everyone. Welcome to our How It Works session on understanding DNS abuse. We have Carlos Alvarez from the office of the CTO presenting for this session. Carlos?

CARLOS ALVAREZ: Thank you so much, Cathy. Thank you, all of you who are here in the room. Also I have a good number of participants online. I understand there are about 23, maybe more as we advance through the session.

We're going to talk about a topic that's very important. It's very relevant. It's also contentious in some respects and something

that people need to pay attention to. We're going to be talking about DNS abuse.

First we're going to lay out some of the sides of the discussion. There are different sides. People understand DNS abuse from different perspectives. We're going to lay out some of those perspectives here, and then we're going to go over some examples of DNA abuse or misuse. We're going to talk a little bit about the evolving Internet landscape at it has to do with DNS. Then we'll end up with a little bit of DNS abuse within the ICANN context.

The first thing that's worthy of mention is that there's no globally unified accepted definition of what DNS abuse means. As this slide that you're seeing says, there are definitional [variants] that include pro-topics such as cybercrime, hacking, and malicious conduct.

You can consider that DNS abuse falls under three categories. It could be data corruption, denial-of-service, and privacy. And there's of course the differentiation between DNS misuse versus DNS abuse. I'm reading a bit what this slide is telling you. Misuse refers to intentionally deceptive, conniving, or unsolicited activities that actively make use of the DNS, or the procedures used to erase domain names or resolve them. We're going to see what this means. We'll see that later on.

What has the GAC said in trying to reach a definition or provide elements to understand what DNS abuse is? What has the GAC said with regards to this broad area? The GAC provided in its communiqué the safeguards that we're applicable to all the new gTLDs in an excerpt of the document that was related to mitigating abusive activity. It mentioned certain abusive activities such as distribution of malware, operation of botnets, phishing, piracy, trademark and copyright infringement, fraudulent or deceptive practices, counterfeiting, and other types of activities that may be contrary to applicable law.

That seems broad, and some are of the view that there are some topics included there that may not necessarily be technical abuse of the DNS as a technical system. But, as I said, we're presenting some of these that the community has. This is one of the views. And there's the other side that doesn't consider trademark and copyright infringement and, in certain cases, fraudulent or deceptive practices technical DNS as such, for many reasons.

There's a big question – I will leave it as an open-ended question during this session – with regards to whether or not spam is or should be considered abuse of the DNS, of the domain name system.

On the side of the operation and security community and law enforcement, spam is seen and considered an indicator and a predecessor of other types of malicious activity. In and of itself, at least until today, it has not been considered technical abuse, as such, of the DNS, the global domain name system.

When you're doing threat research and you're analyzing data and you see that there is a spam campaign that was just launched, you identify the criminal infrastructure that the bad actors are using for that campaign. If you continue to follow their activity, you will sooner or later – usually sooner rather than later – see the follow-up activity that's exactly following up after the spam was sent. It can be malware distribution. It can be distribution of child abuse material. It can be phishing – different types of things. In simpler terms, DNS abuse refers to anything that attacks or abuses the DNS infrastructure. We'll see that in a few slides I have.

There's many ways to see DNS abuse. We're going to talk about two ways so slice this topic. One is the perspective of the abuse of the resolution of the domain names, the technical portion of how domain names are translated into IP addresses. The other perspective that we're going to talk about is the registration the domain names, the probation of registration services.

Those services that are provided by registries and registrars are abused by the criminals in different ways. We're going to talk about that as well.

DNS misuse refers to exploiting the DNS protocol at a more technical level or the registration processes for malicious purposes. All these we're going to exemplify in more detail as long as this works. [I'm aiming] – yes. There.

Sorry. You don't need to read all these slides. That's not the purpose. The purpose is to show you the operational elements of the DNS at a simplified scale.

You have, on top in the blue font word, the authoritative name servers that host the data that's authoritative for each domain name or for each TLD, for that matter. Right below, you have the recursive name resolvers, which you could see as the DNS servers that your ISP – the company that provides you with Internet access – lets you use. They provide you with the DNS resolution service.

Then we have the client or the stub resolvers. What's a stub resolver? It's right here. It's a function within my browser, for example, that looks for the information that it needs to be able to use the resources that I want to use.

For example, if I go to the browser and type www.ICANN.org, the stub resolver function will be called in, and it will resolve the domain name into the IP address, where the content for www.ICANN.org will be hosted. It will download it to my device and I will be able to see it and interact with it, etc.

These three operational elements of the DNS are targets for attacks. We'll see that, basically, everything that's online is a target for an attack.

Examples: this is where it gets interesting. Let's focus on reflection/amplification, which is at the heart of the DDoS attacks (Distributed Denial-of-Service) attacks.

What's reflection? Reflection means that you can send a packet and falsify the information about the source IP address so that the server believes that it was sent by someone else, which will have that server send the response to that other IP address. So if my purpose is to attack Cathy, I will send the packet to the DNS server, and the IP address that I will forward into the packet will say that the packet itself is coming from Cathy, not from me.

If I do that using what are called open resolvers, which are DNS servers that exist out there – there are thousands of them – that don't filter the IP addresses where the queries are coming from and respond to queries from any user from any region in the world, then Cathy is going to be flooded because I will send DNS

queries to those thousands of open resolvers that are out in the wild. All of them will think that all those queries were sent by Cathy, so they will all respond to her. That's reflection.

The other vector is amplification. What does amplification mean? It means that you took those queries that I'm sending – it's tiny. It usually is only a line of command. That line can be as simple as dig name server domain name any. Boom. That's it. That's seven bytes, maybe. It's really tiny. It's just a line of text, while their response will be huge. It can be 2.3, 2.5, or 2.7 megabytes. Multiply that by the thousands of queries that I am having my botnets send. Remember the botnets, those large networks of compromised devices that the criminals operate? They can have hundreds of thousands of compromised devices. The criminal operating the botnet can have all those compromised devices send queries to all those open resolvers that are out there, sending Cathy those responses.

So it's a multiplication effect because of the number of compromised devices that I have recruited and that are part of my botnet and then the thousands of open resolvers that I am leveraging and the way that I send a command – the dig query. Dig is a command that allows you to obtain information from the DNS. It triggers a response. The way they craft that command is such that their response is huge, as I just mentioned.

So Cathy very likely will be knocked offline if she's not on sort of DDoS protection service or what-have-you. If she can't withstand the traffic that's going to come to her, she's going to go offline. There's no way to prevent it.

The first large DDoS attack that [inaudible] that was seen using the DNS as a vector of attack was in 2003 against Spamhaus. Spamhaus is an organization that focuses on not only spam but malware, etc. They do investigations. They provide interesting date for mitigation and protection against threats, etc.

And of course the attacks keep on involving. Now, of course the DNS is not the only attack vector by any means. It's one of them, but as a protocol it is exploited and it's often exploited.

We'll also take about cache poisoning or exhaustion attacks. That's mentioned second to last. The last DNS man-in-the-middle attack we'll also talk about a few slides ahead there.

This is basically what I just described. This is a large DDoS attack that's making use of reflection, sending packets that are forging their source IP address, making all those open resolvers think that it's Cathy's device that's sending the queries. Then the query that they received is such that it triggers a large response. That's the amplification vector. Cathy gets flooded. It's exactly what I just described.

That wasn't me. Talk about going ahead of myself. I had too much coffee. Where were we?

Here. Another way in which you can attack the DNS is by going against someone's name servers. The name server is part of the infrastructure that's used to provide resolution for a domain name. So if I erase your carlos.whatever, I will have to set, ideally at least two servers that the DNS is a global system to to obtain information about the resources that I have associated with that name. In other words, I will define and will make that information available to those name servers – the IP addresses; just to put it very simply, where my mail server will be located, where my web server will be located, where my FTP server will be located, etc.

So if anyone knocks offline my name servers, no one will be able to access that information, which means that I will not receive e-mail nor send it. People will not be able to access my website, etc. So it can have consequences if it's a high-value domain name. Not that carlos.whatever is a high-value target. It probably doesn't exist, but it would be.

The way it works – this type of attack – is that criminals abuse the TCP protocol. When you send a TCP packet to a server, the server responds. This is, again, put very simply. When the server responds to the TCP connection, both the device that started the

connection and the server that's responding establish a handshake, which create a channel of communication that's maintained by both. That means that both have to allocate certain resources to maintain that channel of communication.

So if you have a lot of compromised devices in a botnet and you have them all send responses/queries to a name server in a way that it'll force that server to establish too many TCP handshakes so that it has to allocate resources to maintain those established channels of communications via TCP, you will soon reach a point in which that server will no longer have resources available to allocate any more TCP connections, which will mean that no one else will be able to ask that server for DNS information. It will still be online, but it will not be able to answer any query. As the slide says, name resolution is degraded or interrupted.

If you're lucky, you'll be able to get a response a couple minutes afterwards. Multiply that thousands of times if it's a high-traffic domain name or you may lose resolution. That's the worst-case scenario that everyone wants to avoid.

Poisoning a cache. This is trickery. This is the bad guys being creative, as they always are – well, sometimes. Sometimes they're really done. Remember that we mentioned in a previous slide that, on top, we had the authoritative name servers – like if

I created and registered carlos.whatever and I associated ns1.carlos.whatever and ns2.carlos.whatever to provide the DNS with the information associated with my mail server and my web server, etc.? Those are my authoritative name servers.

Every ISP – many people out there; 9.9.9.9 or 8.8.8.8 or OpenDNS or UltraDNS; there are many DNS provides – is recursive, which means that they ask questions on behalf of someone else. There are some of those recursive resolvers, as they are known, that are not well-protected. They are vulnerable. If you imagine all the thousands of Internet service providers that are out there in all the regions, some are operated by small companies with not too many resources. So they have the infrastructure to operate, but they may not have the resources to protect it.

When those servers are not protected enough, then criminals may compromise them in many different ways, one of which can mean that, if I am a user of an ISP that has a compromised recursive resolver and I send a query looking for data associated with what-have-you – PayPal.com – I may get the correct response, but then since the server is compromised, the criminals will add an extra bit of information there. That extra bit of information there will be something like, "Oh, and by the way, the IP address for BankOfAmerica.com is this."

And it will automatically update my device's temporary memory, its cache memory. In my advice, it would be the host's file. When that happens, guess where my device will take me the next time I want to visit BankOfAmerica.com? If it happens within a defined time period, it'll take me to the IP address that the criminals wanted me to go visit. So that's a bad scenario. That's not nice.

What will happen then? I will visit the criminal's operated server. I will see the content that they want me to see, which in this case would be a phishing site mimicking Bank of America. Then I would gladly provide them with my user name and password, which will not be good for my personal finances, of course.

There are other ways in which criminals can do things like this. They can compromise your device directly and change/modify your DNS configuration. We'll see an example later on of a bad botnet that was taken down four years ago, I think. If my device is configured to send DNS queries to, as an example 1.1.1.1, they change that IP address here or on the laptop, and instead of that legit user-intended IP address, they put their own. They put an IP address of a DNS server that they operate that's configured to provide IP addresses to their own infrastructure. In other words, they will have all the users whose devices are compromised visit their own websites, which, again, is not nice. We'll talk about this in a little bit.

In a case like this, in talking specifically about poisoning a cache the way they do it, they have the user's compromised device send the queries to their DNS server. Say that I'm asking for a site that's not relevant to the criminals – a new site (although anything can be relevant to them). Let's say that news.whatever is not a relevant site to the criminals. Just as this did in the previous example that I gave, they add an extra bit of information to that response that their server is sending my device. That extra bit of information could be similarly the IP address – oh, by the way, the IP address for your bank's website is this. So, again, within a defined time period, I query for my bank's domain name because I want to go into some online banking, and – boom! – I end up at the criminal's website. Again, bad for my personal finances.

DNSChanger is exactly the malware type that I was mentioning. That's exactly what it did. It changed the user-intended DNS configuration. It is pretty pervasive. The criminals that led this botnet – the masterminds behind this operation – were able to make a lot of money. By the time a law enforcement operation took place, law enforcement was able to prove with hard evidence that they had obtained, illegally, $25 million euros. That doesn't mean that they didn't make more. It just means that that was law enforcement was able to prove with actual evidence.

With DNSChanger, the criminals change the DNS configuration on the users' devices. They did something that seemed to be pretty innocuous in that they were replacing the advertisements that users would see when they went to websites. If I went to my preferred news website in the morning at the office with my cup of joe, instead of the legitimate advertisements that I would say, they had replaced those with their own. And it has generated income for them on an ongoing basis. That happened for a very long time. So that's a money machine.

It doesn't harm, or so it seems. Users didn't see any weird behaviors in their devices. They were still able to access the content that they wanted to. They were still able to interact with the Internet and the resources that they wanted to. So apparently nothing weird was going on.

Well, there was. So the action took place, and there was a concern because there were so many infected devices. I don't remember the exact amount of devices, but it was in the hundreds of thousands that were compromised within this botnet by this malware type in a number of countries. I don't want to lie – it might have been around 20 countries. But I'm not super sure.

So the question came when they – by "they," I mean law enforcement – were going to have to do something with those

DNS servers that the criminals were operating. Either they could turn them off, in which case – what do you guys think would have happened if they would have turned those DNS servers off? If all the user devices were sending DNS queries to those servers, what would have happened?

Users would have thought they had lost Internet connection. They would still be connected to the Internet, but their devices wouldn't resolve any name because the DNS servers were off. So they couldn't just turn them off.

Using engineering – to put it that way – they were able to replace those servers. The court assigned an administrator for those servers during a period of time through the [national] certs and different techniques. Awareness campaigns were conducted in all those jurisdictions so that users realized that they had to clean out their devices.

Of course, since then, criminals are criminals. They have found many different ways of abusing the DNS's protocol and the different operational elements of the DNS, some of which are interesting, at least from an academic perspective, because it shows creativity – for bad purposes but it's creative, like covert exfiltration channel. I think this is on the next slide – no. Let's talk about it.

Of course, when you're able to send data from a compromised network without the administrator of the network realizing that they are having their data stolen, that's the covert part. The DNS is considered to be an interesting covert exfiltration channel because usually that little port that's used for DNS communications is not blocked. It cannot be blocked.

Traffic in network travels via ports, from one port to another. The port that's used by the DNS protocol is Port 53. While there are ways in which engineers can reassign that port internally within their own network, it can have some complications. So it's usually never reassigned or moved to a different port. That means that traffic through Port 53 cannot be blocked. It's very complicated to reassign it, so it simply cannot be blocked. If it's blocked, then people will not have DNS resolution, which means that they think that they are fine.

How does it work when you have a covert exfiltration channel? Different ways – at least two that I can think of now. One way, you compromise the device and have that device start sending DNS queries slowly to a criminal name server. The thing is that, within each DNS query, the criminals have replaced the least relevant bits with the bits that correspond to the data that they are exfiltrating. If the engineering team or the administrator of the network is look at those queries, is looking at the traffic, they will still only say that it's DNS queries. They would have to

gather all the DNS queries that are being used for the exfiltration of the specific piece of data. It takes analysis, basically. Realize that the least relevant bits were modified, put together those least relevant bits, try to make something of them, and then realize what it was that was being exfiltrated. That's one way.

Another way in which criminals use the DNS to exfiltrate information, which is a little easier, is through the TXT records. When you create a domain name, you by default will administer or manage what's called a zone file.

In the zone file of your domain name, you define the resources that are associated with it. That's where you include the information for your website, your mail server, and your FTP server. If you have your domain DNSSEC-signed, that information is going to go there. If you have all the techniques for the protection of your customers or the general public – I don't know, maybe some of you may have heard of those. It's more acronyms, unfortunately – SPF, DKIM, and DMARC – which basically simply protects the users, to put it simply.

All those kinds of information go into what is known as TXT records. You can actually put anything in a TXT record. There's no limitation to the kind of text that can go in a TXT record. It's just text. Criminals use those TXT records to exfiltrate information as well. They can send dig queries with information

regarding TXT records out to a name server that then gathers the information and groups them and recreates the data that was exfiltrated and so on.

Fast flux. I think this is mentioned down the road. If not, we'll come back to it, but I think it mentioned.

Domain name registrations are sweet targets for attackers. That's pretty obvious. Criminals and malicious actors unfortunately abuse legitimate providers of domain registration services in both the gTLD and ccTLD spaces. They love to abuse registrars and resellers. They love to obtain large amounts of domain names. It's a very complicated problem to address. Lower prices of domain names tend to lure bad actors, which is just human nature, I think. Cheaper is better, so they just go in and register some more domain names, just as legitimate registrants and users are lured into registering names when they are cheaper. Again, as I said, it's just human nature. There's nothing wrong with that. It's just how it works.

The automated registration of the domain names creates – which, again, is neither good nor bad. It's just how the industry involved. It allows for the operation of large portfolios of legitimate domain names, of course.

But then the criminals had to be there abusing it, unfortunately. When I mention this, I'm thinking of the DGA domains, which is

basically automation in the criminals' hands for the creation and registration or large amounts of domain names.

What's a DGA? It's a Domain Generation Algorithm. Imagine a botnet. Every botnet has to have a command-and-control infrastructure that's used by the criminals for them to be able to precisely command and control their malicious infrastructure.

But what happens if that infrastructure is taken down? Then it will have a Plan B, C, D, E, F, G, and so on. That's what the DGAs are there for. When the botnet realizes that one of those servers associated with command-and-control is down, suspended, or was not defined or whatever, then – boop! – it registers. This is just an example. There are all colors and flavors and variations of DGA behavior. This is just a simplified explanation. Boop! It [registers] any string under whatever TLD, and off it goes.

If the command-control capabilities of the botnet had degraded because of threat mitigation action, then the registration of that new DGA string? Boop! There again. They lost command-and-control capability and just continue their operation.

We're going to mention a very interesting case. I hope we're near that slide. Why do attackers and criminals register domain names for everything? Anything that you guys can think of – for phishing, ransomware, malware distribution, scams, counterfeit

goods, illegal pharmaceuticals – you mention it, they register for it.

The last line – I don't know why it's showing like that – is command-and-control, which is in regards to stability and resiliency. That's the biggest concern because of the size of the attacks.

There are questions sometimes about illegal pharmaceuticals – whether or not that should be considered DNS abuse when it's apparently not related to DNS abuse – technically, at least. It's more similar to the counterfeiting – selling websites, etc. That's true, but there are sometimes things underneath what's just on the surface. I can't go into details, but just keep in mind that there are things that are underneath what you see on the surface. You may see that it's simply a bunch of websites that are used to send illegal, in certain jurisdictions, medicines. But what's underneath it is something that harms a lot of people.

Do you have a question? Would you please come to a microphone?

CATHY PETERSEN: Please feel free to use any of the microphones on the tables. Please give us your name and affiliation if you have any. Thank you.

FARZANEH BADII: My name is Farzaneh Badii. I am the Non-Commercial Stakeholder Group Chair. The question is in my personal capacity. When you say that, under the domain name that sells illegal drugs, there might be some other bad stuff, do you mean there might be some technical abuse? Or are we talking about the content of the website?

CARLOS ALVAREZ: I'm talking about criminal operations that make use of domain names. It's use of the domain names in relation to the content of the websites and more criminal activity that follows after that.

FARZANEH BADII: Just a follow-up. So this has nothing to do with DNS and its technical operation?

CARLOS ALVAREZ: They use the domain name for that purpose.

FARZANEH BADII: Thank you.

CARLOS ALVAREZ: Of course.

So why pay when you can crack? Why would criminals pay for domain names or why do some choose to pay for domain names when they can crack and just gain control of them?

There are different situations in which criminals will choose to, instead of registering, go for the kill and hijack domain names. How do they do this? They can compromise the user credentials. They can compromise the registrants' access credentials to the control panel. The control panel is the web interface that allows registrants to manage their domain name.

Imagine a criminal organization that wants to take over a specific high-value domain name or wants to harm clients of a specific bank. They may just send a pure phishing campaign, a targeted phishing campaign directed at the employees of that bank, after some social engineering, as usual, lure the workers of that bank into a clicking a link that they shouldn't have clicked, and then steal the access credentials.

Whatever happens afterwards goes as far as the criminal wants to. They can simply create a third-level domain underneath the second-level domain that you see. If my bank had suffered a situation like this – let's say my bank is carlosbank.whatever – the criminal can create "I'llphishyou.carlosbank.whatever" and then send e-mail out as part of a phishing campaign, luring

victims more successfully because they will see that the second-level domain is actually my bank's real domain name.

Or they may change the name servers completely. They may change any information that's associated with the domain name. They can change any record. They can take down the entire set of information that's within the zone file for the name.

Then again, there are situations where it has happened, unfortunately, that registrars that may not have their infrastructure appropriately secured have been compromised. It doesn't happen often at all, which is good, but it has happened. When that happens, it's not a good situation. F0rtuatnely, in those very few cases that we've seen, the criminals went after very specific, high-value targets, and the registrars were very quick to respond. That was some time ago, actually. That was handled very appropriately. The criminals after targets that they knew in advance they would be able to gain if they got control of those servers.

So you see it on the user side if the registrants are lured into clicking a link and then the usual phish that gets successful. Or on the attack on the registration infrastructure, if it gets to be successful.

Well, this is another side of phishing. How many registrants could have the same access credentials to the control panel

through which they manage their domain names? How many registrants could have the same access credentials in that control panel as they did in an account that they got compromised in before? In any of the tens of compromises and large breaches that are seen every week, basically, or every month, that's an unknown.

Credential stuffing is criminals attempting to log into as many services as they can with the pairs of user names and passwords that have been compromised in prior data breaches. When they get it, they get in and you're dead. And that's a big unknown. There's no way to test that. That's a big unknown – how many registrants are reusing their passwords for the management of their domain registrations? That's awareness, awareness, awareness, as usual.

Fast flux: it's here. Fast flux is a technique that criminals use that allows them hop from IP address to IP address very rapidly, with the purpose of making the work of law enforcement and threat mitigation professionals more and more complicated.

The way they do this is that they define in their zone files short TTLs. TTL is the Time to Live. That's the time during which the IP address that, for example, may be associated with the website may be valid. After that time, the recursive resolvers out in the wild will know that they have to query again to obtain that

information again. When they do it, they will receive a different IP address. So when you see short TTLs, like 120 seconds, 180 seconds, 2 or 3 minutes, or 4 minutes, that's something that researchers look at and think, "Hmm."

The caveat here is that the large CDNs – the CDNs are the content delivery networks – for their operation to provide stability and to provide load balancing and for other technical reasons also use short TTLs. But that's different. So if you're a threat researcher, you know which are the large networks which use TTLs, and that's fine. But if you come across a new domain that has a short TTL that's associated on top of it with some [newly seen] infrastructure that may have been seen in association with spam or something, then that raises a brow immediately. That's when threat researchers are tempted into blocking traffic associated with that infrastructure for the protection.

What happens when you're law enforcement and you're investigation a criminal infrastructure/operation in which the bad actors are using fast flux, where you see that the content is in this server in this country, and two minutes later, the same content is no longer there but is in this other server in this other country, and two minutes later, the content hops to a different server in another country, and two minutes later the content

hops to the next server in a fourth or fifth country and so on? How does law enforcement address that?

It's hard. Very hard. Double-fast flux – right there. It's a technique that was seen in a huge – how would we put it? – criminal cloud service. It is called Avalanche. Within Avalanche, the bad actors were using double-fast flux. What that meant was that they were changing the name servers very often.

If it wanted to query carlos.whatever, right now, I would query ns1.carlos.whatever. If I queried within two minutes, I would query ns1.cathy.next. Within the next two minutes, I would query ns3.cameron.yoohoo. Every two or three minutes, the name servers would be changing. Then the name servers were on top of that, changing the IP addresses every two or three minutes or the short TTL that the criminals had defined. So that was twice as bad and twice as annoying and twice as complicated, but the good investigators were able to find that out and address it. The bad guys are behind bars. And the head is behind bars.

I mentioned that the DNS is a covert exfiltration channel and that the malware types actually – it's not only for data exfiltration but also and mainly for actual control of the malware that's infected or compromised the devices. Through the DNS, the criminals provide instructions to the devices. The criminals

modify the malware that has compromised the devices. The criminals can inject more malware through the DNS.

It's a headache because, as I said, Port 53, which is use for DNS communications, cannot be blocked. So it's up to the network administrator to have good techniques in place to be able to detect these things.

There are some techniques that I can't refer to right now because that would be another three hours. It's up to them. It's up to each network administrator to implement those techniques.

We just saw this. Two examples of malware types that are doing this, among the many, many more that do it, are Feederbot and Morto. You see here that the botnet command-and-control encodes instructions in DNS TXT responses. So the compromised device sends a query to the name server and the bad guy has configured that DNS server to provide a response in the form – the query was for a TXT record, and that TXT record gives instructions to the compromised device, basically. Those instructions can be anything. Those instructions can be, "Attack these targets and traffic this way." So it can be anything.

The evolving DNS landscape. DDoS a service. Talk about Mirai. Do any of you remember Mirai? Yeah. That is a bad situation there.

Mirai was – how should we put it again? There was an association between providers of what are known as booter services or stresser services with this attack with this botnet.

What's a booter service or a stresser service? It's a website that some kid sets up somewhere, where what they claim is that they sell the capability for you to test how resilient and how stable your servers are. You pay a certain amount of money and what they claim that they provide is, "We will send this amount of traffic during this period of time so that you can test you infrastructure to see if it's resilient and to see if it will withstand an attack."

The thing is, those booter or stresser services sell that service to anyone, whether or not they operate the infrastructure that is to be tested. In other words, they start DDoS-for-hire services. And it's not hard to find them. They can go online and make a very easy lookup using your favorite search engine and you can just find them. There are some that are dumb enough that accept credit card payments, which makes things easier for the light side of the Force. All you need to do is pay and then provide the information and the target that you want to test, because, of course, you want to make sure that it's a resilient network. [inaudible] fine. That's not cool.

They do that via different means, one of which the operation of botnets, of course. We already talked about fast flux and double-fast flux. I mentioned Avalanche. We'll talk about it a couple slides ahead. Avalanche is a very cool case, and you'll see why.

The Internet of Things. I didn't want to mention the v-word that's in between "Internet of" and "Things," but that v-word is unfortunately so true. And everyone knows that it's nothing new.

A good example of how things can go bad: think of the attack – I think it was against Brian Krebs – in October 2016, maybe? September? And against OVH, which by the way is an ICANN-accredited registrar. They're also a large hosting provider in France. They were able to detect that the attack was coming from something around 146,000 digital video cameras.

The botnet had the capacity to send 1.5 terabytes of data. Back then, that was simply unseen. That's insane. That amount of data I can't picture in my head. They got to measure 1.1 terabytes in actual traffic directed against them. It was video cameras. Again, this is nothing new, but it's something worthy of mention, I guess. The DNS was one of the vectors used in that attack – not the only vector, but a vector that was certainly used.

Then we have WannaCry, which is also mentioned ahead, so I'm going to skip it for now.

Avalanche was a cloud criminal service. Imagine you went to a website, created an account, logged in, and you could choose the malware type and the kind of campaign that you wanted to run. What these guys did was run everything for you. All you needed to do was pay them and they would run everything for you. They would provide the malware for you to infect your customers. They would actually infect them for you. They would [inaudible] the domain name for command-and-control on your behalf. They would [inaudible] those domain names for you.

They would provide the hosting for the malware distribution sites. They would operate all that for you. So that was the next level in sophistication with regards to [probation] of criminal services.

Avalanche was of course heavy in DGA registration in domains generated automatically by an algorithm. When the law enforcement action took place, a process within ICANN – that process is called the expedited registries security request – was followed. Through that process, 832,000 domain names were taken away from the hands of the criminals.

So from one instance to the next, thanks to all the cooperation from the law enforcement partners and some folks in the private sector, the criminals completely lost control of their infrastructure. Like poof, it's gone. I mean, it's still there, but

they can't touch it. They can't control is anymore. It feels nice when that happens.

These are some of the strings that were to be created by Avalanche, by the botnet, for its command-and-control purposes. All those 830,000 domain names were to be created under a bunch of TLDs, both ccTLDs and gTLDs. As I said, criminals will abuse whoever they can. They don't care, of course.

Well, there's one thing. Some criminals in certain parts of the world will craft their malware in a way that it will not attack IP address within their own jurisdiction because they don't want their own law enforcement to go after them because it can be really bad for them. So they just skip completely their own IP address space.

The thing, of course, then is that they can't leave their country, which is really not a bad thing. They make prisoners of themselves within their own borders. It's good that they stay there, but it's bad because they do a lot of harm.

This is the outcome of the Avalanche takedown. Thanks to content that is provided by Europol and the FBI – the entire presentation is only to talk about this case. This is just the outcome. Five arrests in four countries, 37 searches in seven countries, 39 servers seized in 13 countries, 221 servers taken

offline, 64 TLDs/832,000 domains in 26 countries, and a lot of victim remediation and awareness raising and prevention. So this is a really huge-scale operation. It is a good thing. It is a big, big, big hit.

WannaCry was a weird thing if you look at from the perspective of the DNS. It was interesting in that, unlike what's usually seen in terms of malware types usually using domain names for command-and-control within any normal TLD – gTLDs or ccTLDs – command-and-control from WannaCry was mostly provided by, I think, seven .onion domains. If you remember .onion, .onion was defined by the IETF as specially-used TLD, which means it will never be in the root, which means ICANN will never have anything to do with .onion. So there was no way to take down the command-and-control infrastructure associated with WannaCry.

However, this young researcher, Marcus Hutchins, the British kid, was analyzing the code. He had managed to obtain a sample of WannaCry. He was analyzing it and came across a string that was within the code, I believe. Of course it was hard-coded in the malware. He checked for it. It wasn't registered. He registered it, and stopped the spread of the malware. By pure chance. He had no idea that was going to happen. Just by registering that domain name, he stopped the spread of the malware.

The reasoning is basically what's up here. If my ransomware does connect to the command-and-control, then it's a process to avoid analysis. Fortunately, that was there. Then the spread of WannaCry stopped.

Then the criminals behind WannaCry attempted registering a second string, but it was also very quickly registered. In the end, the spread was fully stopped. They just moved elsewhere.

DNS abuse is a contentious topic within ICANN. There are different views. There are some who, on the security side, on the law enforcement side, have concerns with regards to WHOIS accuracy and, of course, the impact that the GDPR is going to have on operations and how WHOIS is going to look after May 25th, when the GDPR enters into force.

There are concerns with regards to things like time for response, the time to react, whenever a port of abuse is submitted. There are different types of concerns in this regard.

On the other side, which is also a side that we an organization also have to listen to, is the concern that ICANN should not exit its remit or step out of its scope in the sense that, if it's content, then ICANN should not have anything to do with it. What that translates to is the ICANN contracts not including provisions that would allow for the takedown of, for example, pirated content. That discussion is to be held by the community, not by the

organization. So those are discussions that have to be held by you guys, basically. We facilitate those discussions, but we can't participate in them.

Importantly, the Public Safety Working Group is the home where law enforcement, both civil and criminal, resides within the ICANN structure – the larger sets of ICANN committees, etc. Before the PSWG, as it's known, existed, the law enforcement community hadn't really found a home until, I believe, Beijing, when Laureen Kapin from the U.S. Federal Trade Commission asked Fadi Chehade, the former CEO, whether he would be willing to consider having law enforcement a formal place within the ICANN structure. He threw the ball back at the law enforcement community: "Bring me a proposal." So they did. That proposal is what we know today as the PSWG, which is a working group or a sub-group within the larger Governmental Advisory Committee. That's where they reside.

The purpose of the PSWG is to provide advice to the GAC (Governmental Advisory Committee) and to the larger ICANN community, of course. Some topics that they focus on are DNS abuse, of course, ways in which domain names are used for malicious purposes to harm users, the GDPR, as it's going to have, like it or not, an implication on the WHOIS information that's available for threat research and investigations, and Carrier Grade translation (CGN NAT). In short terms, that's a

technique that some ISPs use. These and other similar techniques by used by some ISPs when they prefer not to migrate to IPv6, basically.

So they instead of having to move up to IPv6, they create huge local area networks and assign their customers internal IP addresses. There are IP addresses that are only meant to be in the public Internet that we would see out, if we were analyzing traffic, and their IP addresses that are only meant to exist within private networks should never be seen out in the public Internet. That's the case, for example, in your company, in your home. Your devices are assigned those private IP addresses.

What these ISPs do is assign those private IP addresses to their customers, even if it's 500, 1,000, or 10,ooo, and they create neighborhood-wide private networks, local area networks, with one single public IP address. That creates a complication for law enforcement because, when law enforcement knocks on the door serving, for example, legal paperwork, or they send a subpoena to the ISP asking for the information regarding the user that sent this traffic from this IP address this day, at this hour, then the ISP will respond:" Well, I don't know. That's 10,000 users behind that public IP address."

And in many countries, there are no obligations, or there may be but they are not enforced with regards to the retention

obligations, maintenance, and storing of traffic logs and log-ins and log-outs. So in many places, you log in, you log out, and poof! The data is gone. No one knows that you were there. The ISP doesn't know you were there. So it's complicated. That's one of the things that the PSWG has discussed in the past. Fast flux, of course, is a technique that criminals use.

These are two simple examples. It's by no means [taxative] or restrictively stuff – contractual probations within ICANN, bigger web of contracts that have to do with anti-abuse. There are many. We could have an hours-long conversation only to talk about anti-abuse from a contractual perspective within ICANN.

I can mention that the registries do have the obligation to monitor their zone for security threats. That means that they have the obligation to look into the domains that exist within them.

If I were the .carlos TLD, I would have to look at all the .carlos domain names and determine which of those are phishing, spam, malware, and command-and-control and report those statistics and metrics to ICANN. That's the obligation on the side of the registries.

Also, if I'm not wrong, I believe the registries also have to provide their abuse point-of-contact information. But I think that's as far

as it goes, specifically with regards to anti-abuse for the registries.

On the side of the registrars, however, it's a little more specific. These more-specific provisions are there in that agreement. That agreement is called the Registrar Accreditation Agreement, or RAA, as we know it informally.

Those more-specific provisions are there as a result of the 12 law enforcement recommendations that were presented by what now is the PSWG – back then it was just the law enforcement community – through the GAC in the Costa Rican meeting of 2012. I think that's when they presented those 12 recommendations.

That triggered the board to order the staff to start negotiations with the Registrar Stakeholder Group. Those negotiations took place during some months, and the outcome is that 2013 RAA that includes a little more specific provisions on anti-abuse.

Some in the operation security community still would like to see clearer and more stringent provisions, but at that point, law enforcement was okay with the text that was agreed to by both the Registrar Stakeholder Group and the ICANN organization.

Some of those obligations, to talk about them quickly, include, for example, that registrars have to take reasonable action when

they receive reports of abuse. Of course, if you ask 10 lawyers what "reasonable" means, you'll have 20 different answers, which makes it complicated. But that's what's there in the RAA.

Another obligation that they have is that they have to provide their abuse point-of-contact as well. That information has to be, I believe, either published on the website and/or on the WHOIS data. It's usually in the WHOIS data, I believe. They also have to publish it on their websites. I'm unsure, as you can tell. I think it has to be there, but I'm not sure.

There's an interesting provision there that's specific to law enforcement. When a law enforcement agency from a registrar's same jurisdiction sends a report of abuse to that registrar – remember, it has to be within the same jurisdiction – the registrar has to provide a human response within 24 hours. That response, as I said, has to be human, not automated. The response doesn't have to be, "We suspended the domain." The response can simply be, "Ack/We acknowledge receipt." That's a valid response.

The person who provides that response, according to the text of the agreement, has to be someone who can basically can decide what's going to happen with the report of abuse; whether or not the domain should be suspended or not.

There are jurisdictions in which this provision helps a lot, where there are many registrars operating from, but there are some jurisdictions in which there are very few, or none at all. So the efficacy or the effects of this provision vary per jurisdiction, of course.

Privacy and proxy providers; if you remember, those services that used by registrants to have someone else's information on the WHOIS output of their domain names instead of their own. If I have a website and I don't want my name to be out there, or my address or my e-mail address, then those privacy and proxy providers that are controlled by registrars also have to provide their own abuse point-of-contact information.

I think that's it. Those are the topics that we wanted to cover. It's a lot. As I said, with DNS abuse, while it seems straightforward in that, when you see a domain name that's used for command-and-control of botnets, it's clear that it is. When you see it, you can do all the technical analysis and there's no way to disprove what the actual technical, scientific evidence shows because it's what's there. But then there are cases in which it's a little more complicated.

So it's a topic that's up for continued discussion. It's up to the community to continue growing and expanding on these topics.

Our job – I failed to say that at the beginning. I am a Director of Security, Stability, and Resiliency Engagement with the Security, Stability, and Resiliency Team. We sit under the Office of the CTO. We do a lot of engagement with the operational and security community and law enforcement.

We seek many things. We try to bring them closer to the ICANN world. We are very interested in having them understand all the discussions that take place here. Just a few weeks ago, a representative from the domain name industry attended a security conference on our invitation. That's the Messaging, Malware, and Mobile Anti-Abuse Working Group. That was Jonathan Frakes, and Executive Director of the Domain Name Association. He had positive interactions.

That's one of the things that we do. We engage. We try to bring people who may have traditionally seen each other as in opposed [benches]. We try to bring them to the others' understandings. If they can understand where each one is coming from, something can built up on that understanding.

We train law enforcement. One of ICANN's duties, as you may remember, is to help maintain the security, stability, and resiliency of the domain name system. That means that law enforcement has to understand what it means when they are looking at a botnet investigation or when they are looking at a

malware distribution investigation. They need to understand how the DNS works. We help them understand that from that perspective so that they can help keep the SSR of the system, as we call it with yet another acronym.

I think that's it. If anyone has questions, please feel free.

CATHY PETERSEN: As a reminder, please give your name and affiliation if you have one.

[MARSY SURMO]: Hello. I'm [Marsy Surmo] from India. [inaudible] question also: whether ICANN has prepared some baseline security standards for DNS implementations or operations. Could it be maintained otherwise? It could be an operational device only and there is no security. All kinds of abuses will happen. Then there will be only post-analyses. So could some minimum baseline security standards be put in place before any operational DNS?

CARLOS ALVAREZ: I would suggest that you look for documents published by DNS-OARC, which is the community of DNS operators. Of course, needless to say, IETF standards that may have some security components with regards to the DNS. Then look for M3AAWG.

About a year-and-a-half ago, they updated what is known as…I forgot the name. If you look up M3AAWG DNS threats, you'll find it. I'm sure. It also provides some good information.

So those are the communities or the groups that I think have worked and delivered documents or standards, as you were mentioning.

[MARSY SURMO]: This is just a guideline. Can we just [inaudible] that before coming on [inaudible] these devices, these minimum security standards are being implemented?

CARLOS ALVAREZ: It cannot be enforced. Anyone can set up, run, and operate a DNS server. Anyone in the world. There's no way to enforce it. It's something that's technically impossible to prevent. There's no rule. There's no binding. Anyone can do it however it was. It's voluntary, which does make it iffy.

Now, with regards to what you were saying, speaking of the voluntary component of it, there are standards and means of doing things that have been defined by the technical community for years; since 1997, for example, the filtering of the [fourth] IP addresses. If you look up BCP 38 or BCP 84, you'll see that those best current practices state 20 years or more. Yet, because it's

voluntary, they have not been widely implemented, or not as widely as you would have expected. It's voluntary.

Any more questions?

Yes, sir?

[HARU AL HASSAN]: Actually –

CARLOS ALVAREZ: Your name and affiliation, please.

[HARU AL HASSAN]: My name is [Haru Al Hassan] from Nigeria. Our challenges in the developing countries is: how do we train our law enforcement agencies to meet the challenges of criminals? Because you have demonstrated a lot of ways in which DNS can be poisoned, can be attacked, how do we train our law enforcement to meet this challenge of the criminals?

CARLOS ALVAREZ: I think a proper route would be to reach out to the ICANN engagement staff in Africa. That would be Pierre. I don't know if you've met him already. Share you concerns with him.

What would happen next is that Pierre, with our SSR team, would coordinate and have law enforcement agencies participate in a DNS abuse training. So my suggestion is to reach out to Pierre because that's a very valid concern.

Yes, sir?

BRENT CAREY: Brent Carey from .nz. I'm wondering if you've got any links into Internet and jurisdiction. Last week I came from Ottawa, and there was a domain name strain. Obviously, infrastructure abuse, registration abuse, and content abuse are all merging. I was just hoping that you've got some links back there, too.

CARLOS ALVAREZ: I don't have them. Go figure. But, yeah, I know that we were trying to organize that forum in Ottawa. Some of my colleagues with ICANN were there.

BRENT CAREY: Because there was a very noticeable absence of the law enforcement area.

CARLOS ALVAREZ: Okay. I didn't know. Maybe that's a conversation with the PSWG. Thank you.

BRENT CAREY: Thank you.

CARLOS ALVAREZ: Okay. Well, there's one more.

UNIDENTIFIED MALE: Just next to it] I [wouldn't] say we do not have a robust mechanism [inaudible] this GDPR for WHOIS. Another way we don't have control over these security issues. It is going to be difficult but [ahead], it looks like.

CARLOS ALVAREZ: What's going to be difficult?

UNIDENTIFIED MALE: On one side, we do not have the authentic WHOIS. With this GDPR, maybe we won't be able to see who is coming to us and where we are going.

CARLOS ALVAREZ: That's right.

UNIDENTIFIED MALE: Secondly, we do not have security of the DNS. We are not controlling it. So there is no control and no authenticity of the person where we are going.

CARLOS ALVAREZ: Let's wait. My suggestion is for you to participate those discussions and provide feedback through the calls from the ICANN organization. It's usually the CEO who's been lately calling for people to provide feedback. So do participate because that's the way in which you can have your voice heard. And it is actually heard. It's not a rhetorical thing to say. It is heard. So throw your concerns in there. That's the right place.

These are some sessions. By no means is this the only session that is relevant to DNS abuse. Just keep this in mind. If you could go back in time, you'd be able to go to yesterday at 11:30 to the PSWG update. Then tomorrow at 8:30 A.M. is the GAC PSWG meeting. I would recommend you attend those two in the morning.

Yeah, I have a slight preference for the GDPR, just because we're at this juncture of things. But both meetings are going to be interesting.

Then it'll be interesting to see what the domain name industry is doing with their healthy domains initiative. It's good to go see

what they're doing because they are doing interesting things as well.

DAAR is a tool that my team developed. It provides information on bad registrations and how they are aggregating more on one part rather than other parts. So that's going to be interesting as well. I'm not going to say much about that because I want you guys to go and attend that session. So just go. Get in there and have fun.

Okay. Well, thank you all so much for being here.

CATHY PETERSEN: Just a reminder: the presentation slides in this session are already in the public schedule. We're going to add the transcripts, as well as the recording for this session also in the public schedule within the next few days. So you can go back and see everything again.

Thank you very much. We're going to have our next How It Works session on Internet Networking at 3:30. Not 3:15, but 3:30. Apologies, as we're going to be running a little bit late for that next How It Works. The session on Internet Networking will be talking a lot about IPv4 and IPv6 protocols.

I hope you can hang around, grab some coffee, and come back. Thank you.

[END OF TRANSCRIPTION]