BARCELONA – How It Works: Understanding DNS Abuse
Monday, October 22, 2018 – 10:30 to 12:00 CEST
ICANN63 | Barcelona, Spain

CATHY PETERSEN: Good morning everyone, welcome to our How It Works Session. For this session we will be talking about Understanding DNS Abuse. Our presenter is Carlos Alvarez from the Office of the Chief Technology Officer at ICANN, he is our SSR Technical Engagement Director.

CARLOS ALVAREZ: Good Morning. As Cathy said, I work with ICANN, specifically within what we call the Security Stability and Resiliency Team. We are a subgroup under the larger Office of the CTO. Our focus is to basically help the Community address issues related precisely to threats that may put the SSR of the system at the risk. We do different things related to anti-abuse. First, for some reason Adobe Connect is not digesting well some of the slides, so I'll just talk to them. It's not many of them that are not showing and the weird thing is that those slides are just text, there's no animations, there's no graphics, its just text and they're not showing. I don't know, the usual glitches that have to happen.

We're going to talk about what DNS Abuse is. DNS Abuse from the ICANN perspective is a very narrow definition if you want. Not that it's absolutely defined, that is still an ongoing conversation but there have been some attempts by the GAC for example and we will get there in a

minute. We'll see what types of activity we consider as DNS Abuse. We tease, as I said, very narrow as compared to what people in the threat research community on in the law enforcement consider DNS Abuse, for us, it's much narrower then for them.

We'll see some examples of DNS Abuse or Misuse, we'll see a bit about the evolving DNS Threat Landscape. Lastly, we'll get to what abuse is within the ICANN context. It's probably good an idea to talk about it right now and then we'll just reiterate when we get there.

Let's go to that last point right now. What's abuse within the ICANN context? The GAC, the Governmental Advisory Committee in its communique of the ICANN meeting in Hyderabad a few years ago, stated that abuse within the DNS, within the ICANN context, should be understood specifically as fishing, moderate distribution and broadnet command control and they left door open for the inclusion of other types of malicious activity. They basically used language sort of saying, "Among others or for example."

I don't remember the exact wording but they left that door open, that's an ongoing conversation. We focus on those types of activities. We focus on threats that may put at risk the stability and resiliency, the operation of the DNS as a technical system. The DNS is needed for users to use the internet as we know it. Without the DNS, users would not be able to interact with the internet, as we always do, as we do every day. They would have to do different things and it would be a little more complicated to interact with the internet without the DNS so we needed to keep it running.

There's no one globally accepted definition. If you ask Threat Researchers, they're going to say that DNS Abuse is any malicious activity that uses domain names or that uses resources associated with a DNS, which are the technical components of the DNS, we'll see that in a minute. That would include for example -- and if you ask an IP Lawyer, someone who's focus and protecting someone's brands or copyrights, then they will say that DNS Abuse involves the violation of trademark rights or copyright rights, that's an ongoing discussion.

We, as the ICANN Organization, as ICANN Staff, we go by what the Community says, so for now we're sticking with moderate distribution broaden the command control and fishing and the conversation continues and it may end up -- it will end up however the Community takes it to. It's up to you all to decided what should be considered DNS Abuse within the larger ICANN context.

From a technical perspective, threats to the DNS fall under three larger categories, data corruption, denial of service and privacy violations. There is a distinction that can be made between DNS Misuse and DNS Abuse. You can think of DNS Abuse when criminals from malicious actor are exploiting someone else's infrastructure, when an attacker compromises the access credentials to a registrants controlled panel and then they are able to create sub domains that are then used for fishing or for moderate distribution, that would be abuse of that person's resources.

Misuse would be if you want the case where the criminals register DJA domains. DJA domains would be domains, algorithmically generated,

that are generated automatically by the code itself of the broad net that the criminals are controlling without human intervention. They just create a mathematical function, let it run and the broad net within certain circumstances, under certain conditions, the broad net without someone there manually typing anything, the broad net itself would create the domains, that would be misuse.

Not that it's a relevant distinction in operational terms but there's an important point there, when you're in investigations you really need to differentiate when a domain is a victim and when a domain is malicious in itself because their response really defers. You can't treat a domain that's being victimized, when for example it's a compromised host and the criminals start using it to launch attacks, you can't treat that domain the same way you would treat a DGA domain. A domain that was created and is used solely for a criminal purpose, they have to be treated differently.

Why is the DNS a target for attacks? A target and a facilitator if you want. First the DNS and going back a little bit on what the DNS is? The DNS is an anacronym that's used for several things, one of which is to talk about the global system that allows for domains to resolve domain resolution is mapping a name that humans can remember into IP address associated with resources, it can be a mail server, it can be a web server, it can be information that's included in the information that's associated with that domain name. The DNS as an acronym is also used -- well, that's that the name of a protocol, a technical protocol that's part of the TCPIP protocols.

People also talk about the DNS to talk about the server that's associated with a specific domain name. If I create Carlos.net and I set a name sever that would be NS@Carlos.net for the resolution of domain, then people who talk about that server as my DNS sever or if you're in a company and that company has a sever that's dedicated to provide their resolution for their internal clients of the organization, then people would also talk about that server as their DNS server. The thing is that on the one hand the protocol is venerable, when the protocols were first created the focus was on making them work. The engineers that were behind the creation of the protocol wanted then to work and security wasn't such a big concern. They couldn't have foreseen what was going to happen later, what's happening today and what's been happening for years.

The criminals of course are really keen and some are really savvy and smart and have figured out some of those vulnerabilities that are within the protocol itself, the technical protocol. Those resources associated with the technical operation of the whole global system, those resources are vulnerable as well, in the end, they're just hardware and software and hardware and software is vulnerable. No matter how much money you put into the protection of an infrastructure it's going to be venerable and it's going to get compromised sooner or later. There are resolvers, servers that are used for the resolution of names, some of which are public, some of which are maintained by small organizations with not too many resources, that are less protected then what you want, then what you would prefer and those are just ripe for a compromise.

The criminals loves them because they compromise those servers, the weaponize them, they use them in different ways attach people from sending data attacks to creating and using them as part of fishing campaigns or spamming campaigns and then even for data exfiltration. The DNS has been used not only once, as part of espionage campaigns. You can use the DNS to exfiltrate data, so if you ran as pie and you compromise your counterparts, your adversaries, you can exfiltrate data from them using the DNS.

Vectors for exploitation of the system. You can initially raise your domain names of course. We'll talk later a bit about this. You can highjack name resolution of registration services and corrupt DNS data. We'll talk about all this, this is all DNS highjack, poisoning, how the criminals abuse the registrars and the registries and how they love discounts. As well, both good guys and bad guys love discounts of course. Discounts bring in a lot of the legitimate clients but the criminals are also going after the discounts because they need cheap infrastructure.

A malicious domain that's used in a spam campaign is going to be -- it's usually listed or included in reputation block lists, about after four minutes since the beginning of the campaign. If I start sending the spam campaign at 9am, it's likely that the domain is going to be seen and included in block lists at 9:04 and within that short time span, the criminals will not only have recovered their investment, but will have made some profit. The cheaper the domains the higher the profit and they just register and register batches of domain names, that it's basically throw away domains, they don't need them again. They

need them to be alive and operational for just a few minutes, that's all they need. From a response standpoint, from a security standpoint, makes it really complicated.

Operational limits for the DNS. What are the components? The technical components that are involved in the technical operation of the DNS? Which are the technical pieces that are there in involved in the resolution of the names? You have on one hand, let's start at the bottom. The client of stub resolvers, what are these? Your devices. I wasn't trying to aim at your head, just the phone. Your cellphones, your laptops, your computers, it's the ones, the devices, the function that send the first query asking for information from the DNS.

When I go to my browser and type WWW.ICANN.ORG, as soon as I click, hit enter, then the browser will call the DNS function and it will trigger a process that will follow. Oh, that's not me. Thank you. As I hit the enter button, that DNS function is called and the query is sent. It is sent to a recursive resolver. A recursive resolver, think of it as the DNS server that your ISP sets and puts their customer's disposal for them to be able to obtain data from the DNS.

That means the ISP has a server configured so that my device can ask quires, send queries and receive responses. The ISP will for me, run all that process. It will tell me where the website for ICANN.ORG is located, for WWW.ICANN.ORG is located, it will obtain that IP address so that my browser can connect to the server that hosts the content and downloaded, so that I can see it and interact with it. Those are resolvers.

The resolvers have a feature that allows them to store in their cache memory, those questions and responses that they ask and receive. The server of ISP or my organization's internal DNS server, they have that feature. They ask the questions, they receive the response and they store that data temporarily for a timeframe that's called the TTL, that can be a shorter or longer timeframe. We could talk about only TTL's for so long. Then the stub resolvers also have a temporary memory that's specifically DNS data. In the hosts file in new devices, you would see all the associations between domain names and responses that have been asked and received. Remember, DNS poisoning, we'll get there later but remember that these two, they're recursive resolvers and the stub resolvers or the clients, they have that little temporary there and that's usually a target for attackers, changing what's in that memory, that's poisoning.

Then there are the authoritative name servers, those are the servers that have the answers. The answers with authority. That means for example if you're a police investigator, if you're doing a low enforcement investigation and you need to provided evidence within a criminal procedure, you'll have to be able to ask a name server with authority, in order for you to be able to provide that evidence within the criminal procedure and having that evidence come from that authoritative source.

If you obtain, even it's the same information, from a server that doesn't have authority, that's not the authoritative server for the domain that you're looking in to, then it could happen that the counterpart of the judge himself or herself would question the validity

of the information that you provided in the criminal procedure. It's just knowing who to ask, what server to send the queries to. These have the authoritative name server. The authoritative name server in the example that I gave a few minutes ago, if create Carlos.net and I associate NS.Carlos.net or .com, so that it's not a bailiwick, if I associate that name server to my domain, that's going to be the authoritative name server and that's where the whole DNS is going to find the answers for my domain. That's the server that's going to tell the whole world, where's the IP address? What's the IP address for my mail sever? For my webserver? That's where I can include the DNS sec records, etcetera.

Of course, this is not showing. Adobe Connect, not that we can help it, it's not digesting well everything. The intention of showing this was basically to show the landscape, the threat landscape on the DNS, from the technical point of view. These are the authoritative name servers, we have resolvers and the stub resolvers and here you would see what the attackers would be targeting in each of these components of the DNS from the technical resolution perspective. So, to say, the landscape is pretty wide and they attack and they target different aspects of the operation of the components of the DNS.

Some examples of DNS Abuse. This is just a list of what we'll see next. Data attacks. The first large attack that made use of the DNS was an attack vector, was as far as I know, in 2013, March I think, against Spam House. A group of attackers that had been ousted by Spam House, Spam House is an organization that does a lot of research, of course on Spam but then also on modeware, botnets, fishing and

other types of malicious activity. They ousted an individual who they thought was involved in the operation of a large spam operation. This individual got pretty upset and decided to retaliate with his friends and they launched a DDoS attack, a distributed denial of service attack against the name servers used by SpamHouse.org.

That was a well thought attack because they were able to disrupt SpamHouse.org. They had to get on Cloud Fare and they were successfully able to mitigate the attack but the immediate result was that SpamHouse.org was down for a couple hours. The traffic associated with the attack reached 330 gigabits, which is a lot back in 2013, that was much more then it is today.

How do they do it? You are the attacker. You are in control of a botnet. Remember, a botnet is a large network of compromised devices that thanks to the modware that you've implanted or that you've used to compromise those devices, you can ask those devices to do many things and you can do many things with those devices. One of those things is use them as weapons against someone else. On the one hand they have that large army of zombies that are waiting for the botnet controls orders. Then the criminals identify DNS resolvers, remember they're recursive resolvers that I mentioned, they identify those resolvers that respond to queries sent by any user.

There are different ways of managing or running a DNS server. You can, if you're an ISP, you can run your sever in a way in which only your clients, those who have IP addressing assigned by you, will receive responses. Let's say my IP range goes from 1.1.1.1 to

1.1.255.255, any user within that IP range will receive a response from my DNS server. Any user outside of that IP range will not get any response. That would be one scenario.

Another scenario is one in which it doesn't matter what IP address the DNS query is sent from, your server will send the answers, that in and of itself is not an issue, those are known as open resolvers because, as the name implies, they are open to answering DNS queries sent by anyone, regardless of whether or not they are customers or regardless of whether or not they have any relationship with the organization that runs the DNS server. The key thing is what's called rate limiting policies.

Think of 8.8.8.8 which is Google's or 9.9.9.9 or 1.1.1.1, there are some public resolvers that are very well managed, that grate limit the traffic and what grate limiting means is that they will not send responses for more than -- think of any number, more than 10 queries per IP address per hour and then they will stop responding or 100 queries per IP address per hour. Whatever the operator the open resolver wants to implement as rate limiting policy.

Unfortunately, there are many, many, many according to Shadow Server. According to the Shadow Server Foundation they are continuously scanning the entire IP for space to find open resolvers among other things. According to them there are millions of open resolvers out there that are not rate limiting. There are millions of servers that the criminals can use to send queries. The way they do it, let's say I'm going to attack Cathy and Cathy's IP address is 3.3.3.3.

You send the IP packet, you send the DNS query and forge the source IP address. It's very easy to forge traffic that's going on the wire.

The information that you're sending, you can forge, you can go into the structure of the IP packet and modify the information that's within it and you can tamper with field that specifies which is the IP address where the DNS is query is coming from and instead of leaving the original IP address, you can include your victims IP address. With what purpose? So, that the DNS resolver that's receiving the query, thinks that it's Cathy who's sending the queries.

Let's see the whole picture. You have the botnet control who has thousands of compromise devices at his or her disposal, then you have millions of DNS resolvers that are responding queries to anyone. You have those thousands of compromised devices that are sending queries to the millions of DNS resolvers, in every query telling the servers that it's Cathy who's asking. What are those resolvers going to think? That they have to answer to Cathy and they are so kind and so efficient that they will all respond to Cathy.

Cathy will get thousands of compromised devices times of thousands or tens of thousands of DNS resolvers, which is the application factor, the reflection is being able to tamper source IP address and then they command that they send to obtain the data from the DNS in a way in which the response that each query triggers is large, it can be 2.5, 2.7 megs depending on how they do it and depending on the information that's in the DNS for the domains. Cathy is going to be in a really rough situation, she going to get flooded, she's going to get tons and

tons of data and if she's not prepared to receive that attack, she's just going to go offline.

There's no way in which she can survive and attack like that. Simply because of the amount of the traffic, it's insane. Remember, five years ago, almost five and a half, with this attack against Spam House, it was 330 gigabits, that number has only gone up. There's just so much more bandwidth today. Of course, the DNS is not the only attack vector that bad guys use, they use NTP and other protocols as part of the attacks but we're only talking about the DNS because the second and that's our focus, not NTP or other stuff.

That's one way in which the criminals can use the DNS. Another way in which they can use the DNS as an attach vector is that the order of the compromise devices that are part of the botnet to send DNS queries to the target name server, they want to take down Carlos, NS.Carlos.com because they want the resolution for my domain name to stop. Why do they want that? Because if they maintain that attack for long enough, no one is going to be able to access my website. No one is able to send me email. No one is going to be able to connect to my FTP serve.

What they do is that they ask the thousands of compromise devices that are within their botnet to send queries to again, open resolvers, the ones we just mentioned and in turn those open resolvers will establish TCP connections with the name server that's going to be their target with the victim name server. The fact that its TCP connections means that the name server has to allocate resources for

the maintenance of the TCP connection. The TCP connection, you can think of it as a three-step handshake. There's what's called a Sin, which is like," Hey buddy, how are you?"

Then the server responds, [inaudible] which is, "Hey, I'm good, what's up?" And then the server or the device that initiated the connection responds, "Cool, I'm here, so let's talk." That's the ACC [inaudible] ACC, and then the TCP connection is established. Once it's established because it's ongoing and it's maintained over time between the two, the device and the server, then the name server has to allocate resources and if you establish too many TCP connections with that name server, then it's going to run out of resources at some point. Once it runs out of resources, it can't establish more TCP connections, which will mean that it won't be able to answer any questions. It won't be able to receive queries and it won't be able to answer. No one will be able obtain information from that name server and if it's an important name or an important date, then off it goes and no one's going to be able interact with it.

Think of a website used by a company who sells tickets for shows and games and things and three days right before the game or the big Christmas show the criminals are able to launch a long enough successful attack against the name server for that website. People are not going to be able to purchase tickets, they're not going to be able to make reservations, they're not going to be able to send emails to that company because it's going to be -- the resources for that name server are going to be depleted. It can be pretty awful. Name resolution is degraded or interrupted.

Let's continue. Poisoning a cache. There are two ways in which the attackers can compromise the cache memory. Remember I mentioned how the stub resolvers, which are our devices, they have these host files that includes the association between the domains that have been asked for recently within the TTL, the time to leave and the responses that were received from the DNS. The host file lists domain names and IP addresses pretty much. Then the recursive resolvers, which are the DNS servers operated by the ISP's, they have the cache memory in which they keep the same information.

Then there are routers at the organizational level or home routers that also have DNS functions, that also have a cache, a DNS cache. Those are targets for the criminals. Why? Because the criminals want to be able to change the IP addresses associated with domain names. Imagine what happens if you are an ISP and provide an activity to 20,000 customers, all of the use, it's just a hypothetical example for the sake of the explanation, all of those 20,000 customers use one DNS sever and of those 20,000 customers, a large proportion of them would use banka.com for their financial servers, that's their bank and they go to banka.com every day to check their balance, see their statement, make payments, etcetera.

Well, it also happens that one day the criminals are able to get into the ISP's DNS server and change the IP address for banka.com. The result is that the next time the users are querying that server to obtain that IP address, they will get an IP address and their browsers will take them somewhere but that somewhere will be a server, a web server operated by the criminals, which is nasty, which can be really bad.

When it happens at the level of the ISP, people know it as farming because you're basically just farming on a lot of different customers.

Imagine, that could be 15,000, 10,000, who knows how many users that are at risk because of that slight change in the temporary memory, in the cache memory of the DNS resolver used by ISP. When the browser resolves and when the browser downloads the content and renders and the user sees the website, it's likely going to be an almost real time copy of the legitimate website. Users shouldn't be expected, they should be using the latest versions of the browsers, they should be able to recognize when the digital certificates are missing but in reality, we from the security side, our job should not be to expect them to be on the latest software. We shouldn't expect them to realize when a digital certificate is missing.

Remember, the little lock that shows when you connect to a https website. This is to say that many users will just go in, they won't realize that they are connecting to a server that's not their banks and they will provide their user name and password and bye, bye money. The same can happen at the router level or at the device level. That information will be changed.

Think of if your company is designing intellectual property, if you company owns intellectual property assets that are very valuable, you've just about, within a few months to release and important invention that's going to make your company because it's a great, really cool invention, then it can happen and it does happen, that certain adversaries will compromise the DNS, either the resolver or the

router, depending on the network configuration, within your organization, to be able to see what they can steal from you because whomever owns the DNS traffic at a target, at a big time organization basically owns the victims traffic.

It's really bad. Also, this touches a little bit on DNS hijacking. DNS hijacking is a little different, we'll see it here but let's just mention this right now because it's up here. It's a combination of hijacking and poisoning. Say, your device like my little laptop there, it's configured to send the DNS queries to a given DNS resolver. The modware compromises my device and it is able to change that IP address of the DNS server that I want to send the queries to.

Basically, the criminals would point my device to send the DNS queries to their own DNS server, that's what this is. It's what happened with DNS changer. DNS changer did exactly that. If my DNS queries were going to 9.9.9.9 they would change it and instead my DNS queries would go to whatever, 6.6.6.6 and they would send me the answers that they're wanting to, whatever they wanted to.

Again, for stealing user credentials to access banking services, change publicity ads; the criminals running DNS changer made a lot of money just by replacing at's, which is the less harmful thing, not that it's legitimate, not that it's legal but at least it's not infecting malware. They also did many other, way more malicious things but simply by replacing the at's, they were able to make a lot of money.

The connection between hijacking and poisoning is that if you are operating a name server or recursive server, I keep confusing them, you can send the replies that you want to the devices that are asking for questions. If I send a DNS query to a DNS server that's operated by the criminals or the ISP legitimate DNS server but it got compromised by the criminals, if I ask for let's say WWW.ICANN.ORG, then that server that they criminals operate or that they compromised will for sure it will reply saying, the IP address for WWW.ICANN.ORG is this but then also they IP address people that come is this and that information is going to get stored in the cache memory, it's going to get stored in the host's file so that my device does take to me WWW.ICANN.ORG but then the next time that I go to PayPal.com, then it will take me to their server.

I wasn't even asking for PayPal.com but the criminals are so proactive and so efficient and so kind that they gave me the information even when I wasn't asking for it. They'll steal my PayPal account because I'll go in and very voluntarily I will be providing them with my user name and password. This is DNS changer.

Now, DNS protocol and registration system misuse, domain name registration hijacking, DNS hijacking, let's see some examples. The process for registering or creating domain names is across many registrars until this is automated. You don't need to have someone there typing things in a keyboard for the domains to get created. You can find inexpensive domains at prices that are really, really low, and that of course as I was saying almost at the beginning, that's good for the legitimate customers who want cheap domains, that's legitimate

but also the bad guys are going after cheap domains. The cheaper the domain they hire, the higher their return investment, and they are always following -- they're always keeping an eye on their domain and market to see where the cheapest prices are, and they go in and they raise large batches of domains that are just throwaway domains.

Now the criminals use domain names for the operation of their infrastructures, it gives them an extra layer of flexibility. Before for spamming they used mostly only IP addresses and the IP addresses were sort of hardcoded in their malware that was used to compromise the devices that were then turned into spamming servers, internal servers that would send out the spam. But those IP addresses are very easily identifiable. Once you see that IP address then it goes into a blocked list and [inaudible] because they know providers and the ISP's are blocking that traffic.

So with the domain names the criminals have a little more flexibility, they can change their IP addresses, they can changed named servers which is what they were doing with Andromeda and with Avalanche, we'll get there in a minute. So that's why they love domain names because it gives them a lot of flexibility.

And then if you as I said earlier as well, if you're operating a botnet you need to have control of that botnet. The way they keep that control of those resources of those devices that are out there in many countries is by being able to communicate with those devices, and for that you need the device to be able to call you. So the first thing that happens when a device gets compromised is that it calls its dad.

It gets compromised, it may do some working scans within the internal network and then it calls, “Hey dad I’m here,” which is telling the botnet controller, “You have a resource it’s available, this is the malware that was installed there, what do you want to do now?” And then the criminal can do whatever they want. They can extract data, they can weaponize that compromised device to send a text, et cetera, et cetera. But all facilitated by the use of a domain name, that’s the command and control server. Yeah, so yeah.

The automation and the [inaudible], the domains gets created and of course makes things easier for the legitimate industry. That’s how it should be people -- legitimate users with legitimate interests make use of those services all the time and that’s okay, that’s legitimate obviously, but then the criminals abuse those services. And that’s hard to tackle. It’s not easy to be able to stop the abuse of those automated systems. Why do I talk about criminals raised for domain names, well I mentioned the flexibility that domain names add to criminal activity and there’s everything if you look at malicious activity.

I don’t how risky it will be to say that all types of criminal activity make use of domain names, most criminal activity does for sure make use of domain names for phishing. You’ll see of course all the types of look-alike domains that are created by the phishing gangs that are quite creative. They not only create the domain that can look very similar to how the legitimate domain looks like, but they will also create a digital certificate associated with that phishing domain so that when you go to the phishing website you’ll see the look. And Chrome and Firefox,

those latest versions will let you go through because there's a digital certificate, and it will be encrypted from your device to the web server. Yeah, it's encrypted, there's a digital certificate but it's a phishing website.

And they can do more things. You can include some records, some information in the DNS of that malicious domain to make it appear more legitimate. So it can be pretty nasty for an unaware user because the level of sophistication that the bad guys are reaching every day more is making things more and more complicated for my mom, for my sisters, for my aunt. Because they want to be able to recognize such sophisticated types of criminal activity. Ransomware payment of webpages of course.

Speaking of ransomware, there was this interesting case that probably you all remember, WannaCry It's mentioned later in some slides, but what I wanted to say right now is I think one time where WannaCry was ransomware where I was a big partner, but one of those two I don't remember which one was using domains in the gTLD space but then most of the commander control mains were Dot Onion Do you guys remember Dot Onion? Yeah, some yes some don't. Dot Onion is the extension under which the hidden services exist within Thor, Thor being this anonymization cryptography based anonymization -- whatever.

Dot Onion has this particular in a way because the IETF, the Internet Engineering Task Force defined some time ago that Dot Onion is a special use domain, which means that it will never be on the route.

It's never going to be in the route, and ICANN's remit is particularly specifically related to the route, so we'll never have anything to do with Dot Onion aside from not including in the route. Fortunately we haven't seen that many botnets or criminal infrastructures making use of Dot Onion domains. I'm sure there are of course but not big bad wolves out there harming lots of people, so that's positive. Because when you're dealing with Dot Onion stuff there are added layers of complication, serious complication.

Malware distribution sites, scam sites kind of a good site to lead a pharmaceutical sites. There's everything. And as I said, depending on who you ask, you'll receive different responses with regards to what could or should be considered DNS Abuse. For us it's narrow, we'll gather again and I'll repeat what I said earlier. One thing that I'd like to mention right now is that because the ICANN scope is so specific, because what we can address is so narrow, while there are some provisions in the contract and the registrars the companies that offer the domains to the public, provisions specifically related to internet abuse, and I'm thinking particularly it's the 3.18.

It says that registrars have to take reasonable and prompt action when they receive reports of abuse. However, that provision should not be used for trademark infringement, corporate infringement, freedom of expression. It's not up to the registrars, but they will always respond when someone be it a trademark owner, or corporate owner, law enforcement officer pursuing the takedown of a domain involved in trademark infringement, or --

There was a case where law enforcement agency was hoping to get a domain suspended because it was posting information, basically encouraging teenagers to commit suicide. The registrar replied saying, "Nope, I'm sorry. We're an American company, we're based in the United States, there's the First Amendment, we cannot suspend." So that poisoned the appropriate path for that law enforcement agency to pursue. In those cases the registrar will always require a court order from the law jurisdiction to.

When we interact with law enforcement or threat researchers we'll always clarify that. If you're pursuing piracy or trademark violation of things that are sort of First Amendment like freedom of speech, go to the court, don't even waste your time, don't make the registrar waste their time. We said that's not the path.

Illegal pharmaceutical or piracy sites. Usually illegal pharmaceutical can be a controversial topic. What you see on there on the surface is that those are services that allow people who need medicine to obtain that medicine at a lower price, at a cheaper price. It'll be illegal depending on their jurisdiction, it'll be illegal depending on whether the medicines or the -- yeah the medicines are being sold to.

But the point that I want to get to and I can speak about because it became public a few years ago is that, there was this one case where a company was focused on the registration of illegal pharmaceutical domain names. They always had between five and seven thousand domains, and all the content in the websites was a replication of about seven templates. Really poor quality, low resolution pictures,

always the same nurses and the same doctors, and the same stethoscopes, the same fonts, the same color but really poor quality, but they did get to make a lot of money.

Hide post concerning. Well yes, that's illegal activity in some jurisdictions, but they were using that money to buy weapons that they then sold to an organization that's considered by the United States and some other jurisdictions as a terrorist organization. So you see on the surface one thing, but underneath it there are things that can be much more complicated. Yeah, so there are interesting things when you look into a little more detail.

So there are seats here if you want. Okay. And also we mentioned the criminals operate named servers for the resolution of their operations of an infrastructure like in DNS Changer they have their own DNS server, and they need names for commander control to maintain their control of their botnet, of all the compromised devices that they're using.

Something that's very, very nasty and that can harm a company not only in the wallet, but also reputation wise is when … Imagine this scenario, there's a criminal gang they want to abuse a domain that has very high reputation. I'm not going to mention anyone, let's say an example that comes that's near enough. Example that comes is a worldwide known brand with millions and millions of customers. They send a technical team that's responsible for the operation, for example that comes from phishing e-mails, and then unfortunately one of the engineers clicks on the link, damn him he shouldn't have

but he did. The result is he receives a phishing e-mail claiming to come from the registrar, with which for example it comes registered. So he goes into this phishing page that looks like the registrars control panel, he provides his username and password to the control panel and as a result the criminals are in. They now have access to the entire administrative panel of example that come.

What happens next? Basically anything. Whatever the criminals want, they can change all the WHOIS information, they can change the account information, they can change payment methods, they can include sub-domains. Sub-domains are third-level names under example that come, they could create -- and what they do sometimes is they create sub-domains depending on what they're looking for. They can create sub-domains. Like one used sub-domains.

So the sub-domain gets created, any victim in the world interacts once with that sub-domain it immediately disappears and new ones get created instead to make their investigative work of course much more complicated. Or they can just set those sub-domains to host phishing pages, or malware distribution points uploading the malicious pay load on them. But the thing is that the company that runs example that come, the super famous brand worldwide has no idea what's going on. Why, because there website works, e-mail works, operations continue just as normal.

So if the technical team is really not onto it, if they're not really paying attention they won't realize until it's too late. And it'll be when they realize it's probably because their domain was included in their

reputation block list, and if so they won't be able to send or receive e-mails, traffic will be blocked by many ISP's coming onto from their domain. So at that point it's just an embarrassment, because everyone knows, "Whoops, you guys got compromised," and with all the consequences that may come from that. And then they have to go in and clean it up, and touch base with their reputation blocked list providers and tell them, "Hey, clean it up please, remove this," and the whole thing that can be a pain.

This is basically what these slides. Same thing criminals destroying registrar e-mail correspondence that's in somewhat often -- there was this attack group, lizard squad that some time ago was really keen and really interested in stealing user credentials to administration control panels. To the registrants they wanted to steal registrant username and passwords to be able to do things with the domain names. And they targeted some high value names which I will not name. Then what ended up happening with that criminal gang was that their, let's call him their leader, was arrested in Finland.

There was a very cool investigation and operation, he was a minor as it sometimes happens, and then the defense attorney argued within the proceedings that he suffers from Asperger Syndrome, so he was released. I can't comment on whether or not that's real, if he does then by all means hopefully all the assistance and the help that he needs. But if that's not true then that's just truly unfortunate, because he's out again just doing whatever he wants.

Fast locks; fast locks is a technique that allows the criminals to be hopping from server to server making their investigations harder. Remember that I talked about the TTL? TTL is the time during which the information associated with a domain name is going to be valid for. In the information associated with my domain name I can define that the IP addresses are going to be valid for one minute, two minutes, three minutes, five minutes measured in seconds, and that's in the zone file of the domain name. That's a value there right next to the -- it's called the start up authority. Assured TTL to some in the threat research community it's indicative of a red flag. A red flag simply means that it's one of the many indicators that may suggest that the domain is malicious. Not because there's a small TTL and short TTL, it won't necessarily mean that the domain is malicious but it is a red flag.

So they go in and look into it a little more detail to make sure whether or not there's something else if the domain is or not malicious. And why is that the case? Because if I run a malware distribution website and I define a short TTL, let's say 120 seconds, two minutes, there's a law enforcement agent who's investigating what I'm doing, they are able to find the URL's of the points where the payload is available for the victims to download and get compromised.

Whenever they get to those to download for analysis in their sandboxes, the content that's associated with those URL's it's going to be already too late because it's really hard for them to be able to access those URL's and download the malicious content while it's still there. So that content is going to be jumping from IP address to IP

address and that means that malware payload is going to be jumping from country to country across many countries depending on the availability of IP addresses that the criminals may have. So it makes things complicated.

If you are the law enforcement investigator or you're a threat researcher and you have to analyze the malware sample that you received information about, it's most likely that you're going to get there late. You'll try to download the sample, but when you connect to that server it will have already moved once, or twice, or three, or four times. So it's complicated, it's really complicated. It's called fast locks.

And then double fast locks, because of course the criminals are not happy with just adding one layer of complication they have to add more, then -- So that simple fast locks is defining the short TTL. The short TTL remember means that the IP addresses will be valid for just this short time, two minutes, three minutes. With the double fast locks what they do is before that the name servers will be changing every two or three minutes, or five minutes.

So if a criminal and I operate Carlos.net not only the IP addresses associated with my domain will be changing every five or so, but the named servers themselves. So right now the named server is going to be enerzedcarlos.net, in five minutes it's going to be enez.whathaveyou.xyz, in five it's going to be enez.cocacola.whatever,

in five minutes it's going to be a different one. And then you query those named servers and the IP addresses they provide will be changing also every five minutes. So if you're an investigator you may as well just bang your head against a wall basically, because it's really hard to investigate. You have to do different things then this on your investigations.

Okay. I have activity in the back. As long as those kids are for us I guess that's just white bag lunches, now brown bag lunches. Anyway, that is the case with Avalanche, Avalanche was a big nasty platform that was taken down in December of 2016. Worthy of mention is the fact that during the operation law enforcement made use of a process that exist within ICANN that's called Expedited Security Request. Of course because we're all about acronyms, whether we like it or not, that process is -- we simply call it the RSR, and it's been quite successful fortunately.

Through the RSR what happens is that with court orders and all the paperwork as usual, the domains that currently exist that are currently being used for commander control of the botnet as well as the domains that will be raised within the future, that the algorithm will create in the future, are taken away from the hands of the criminals. And with that they lose complete control of the infrastructure. All of it goes, "Boom off", goes. They lose it, which I can picture them at their dark caves or bright offices because there's everything and then, "What the hell, what happened, everything's off. Dude did you do something wrong? Did you get the domains?" "Yeah I did." And then they don't understand, and then they realize, "Oh shit."

For Avalanche the bad guys didn't really have that much time to react because as the RSR was implemented by many TLDs, it was 64 TLDs, an elite unit of the Ukrainian police was not knocking but banging, throwing down the door of the guy that was leading that operation. His name is public so I can mention it publicly because it's already public. His name is Gennadiy Kapkanov, and yeah that is a big. Yeah, that was him.

And the thing about Avalanche is that it was the next level in crime … cyber crime sophistication. It started as a moderate distribution service, but then it evolved over time and ended up becoming a full-blown service on the cloud where they would offer their clients the registration of the domain names, the hosting of the malicious content where their victims would get compromised with a Trojan malware. They would offer layers of protection so that they infrastructure wouldn't be easily detected by law enforcement. They were reall well thought off, they were smart; not smart enough fortunately, but that investigation took several years. Like seriously several years. It's probably the longest investigation that I've known of -- that we've know of in relation to something that has to do with abuse of a DNS. It is shut down so off they go.

The DNS is a covert infiltration channel. I mentioned a little bit after the beginning how the DNS -- How much time do we have left Cathy? Like 20? Great thank you. How the DNS is used for espionage campaigns and these are some of the clever ways in which the criminals abuse or exploit the design of the protocol and the operation of the technical components that are part of the resolution of the

domain names. There are different ways in which you can exfiltrate data from a compromised device. I guess there are at least two ways, you could of course find others but let's talk about these two ways in which you can exfiltrate data by sending DNS queries.

That's the whole thing, you convert each DNS query that the compromised device is sending into a box that can change the data that is being exfiltrated. So it's still a DNS query, it's still asking for a domain name, but within it you're sending something. So within every IP packet there's a field that's called the padding. The padding exists so that the size of the packets complies with a certain given technical standard. They have to be of certain size when they are on the wire.

Let's say the malware has to exfiltrate all the word documents within the compromised device, the malware compromises the device, identifies all the target documents, calculates the size of those documents and calculates how many DNS queries it's going to have to send out to be able to exfiltrate those documents. It's a very simple calculation. If I have this many DNS queries that means I will have these many zero's and one's in the padding, so I just have to send as many DNS queries until I put all the zero's and one's of the documents into the available zero's and one's of the padding, and [wheet] off they go.

And they send out the queries in a way that will not trigger the alarm. If there's a spike in DNS traffic all of a sudden, if suddenly the amount of queries being sent through the port 53, which is the tailor port that the DNS uses. If there's a spike and all of a sudden you see 10 000

queries then probably the network administrator is going to go in and see what's going on. But if there's just queries coming out every 90 seconds, every two minutes that's not going to trigger any alarms. And if there's no DNS traffic monitoring which unfortunately not many organizations do, no-one's going to notice. It's just going to go [whistle], out slowly.

In an attack like that all the attacker needs is to get the information out. They're not in a rush, they don't care, they can wait, they're not running against the clock so they just sit and wait. The DNS server that's receiving those queries is configured with a specific [inaudible] with the zero's and one's identify the beginning of each file, the file type, the flag that marks when the file ends and then reconstruct the files that are being exfiltrated and put them together. And the same can happen -- Every IP packet has what are called least relevant bits, you can change the least -- Yeah, least relevant bits; you can replace them with whatever zero's and one's that you want and the packet will still be the same.

It's still going to be a DNS query, it's still going to be asking for a domain name. So you can use the padding with the least relevant bits and send out the information that you want, all you need is to have the named server on the other end configured to identify the files, the beginning, the end, the type and just put them back together and off you go.

It's mostly never detected. Another caveat is that traffic on port 53 as I said, is a [inaudible] port, I'm sorry, that's used by the DNS. Traffic in

that port is not encrypted, and shouldn't be encrypted. The criminals don't need to encrypt the traffic that they're sending out when they exfiltrate data, it's just DNS queries. If they encrypted it then that would be an anomaly and that would likely make the network administrator curious and he would go in and say, "Hmm what's going on?"

But this is not encrypted, it's in plain view. And say their network administrator gets curious because he's seeing every 90 seconds always there are queries coming out from these two devices, he goes in and sees some of those queries, it's just queries. And the queries will be something like big@nameserver domain name, something a or n or mx. That's the composition of the query. That's it, nothing weird. So unless he was able to see all the queries, and put them together, and do a lot of analysis he wouldn't be able to tell.

Another way in which the DNS is used, the other way around to receive information is when the device gets compromised with the malware, sends a query asking for the -- connecting to the commander control server and through a response to a DNS response, the commander control server injects more malware. So in this case a DNS query would be asking for what's called the text records. Text records are information that you can associate with a domain name, and since as the name suggests you can include anything, it's just text.

You can literally write a love letter to your girlfriend and put it in the text record of your domain name, and it'll get all of it. So the commander -- The compromised device sends a query, the

commander control server receives the query and sees that it's asking for a text record. What's included in that text recorded? Encoded malware, and it sends it back, the malware in the compromised device recognizes that the text record is the encoded malware, decodes it and further goes the compromise, all the DNS.

Okay, well there's IoT stuff, just to add more to this whole thing. I talked about Avalanche, a huge super sophisticated -- sophisticated, not super sophisticated crime platform. The IoT stuff it's botnet equipment to the next level.

In, when was it, September of 2016 Mirai, these big bad botnet launched attacks against Brian Cripps dying DNS, there you go, and the French registrar and hosting company, a large hosting DNS and registrar company OVH also got hit with some of that traffic. OVH did a measurement and they published -- their CEO published it on Twitter as they were finding out information. What they were able to identify was that the botnet was comprised of 147 000 CCTV cameras. All video cameras, IoT stuff, and all weaponized launch the attack against them. The attack was not using the DNS of course there were several other vectors in the attack, but it was using the DNS to send traffic through.

They measured 1.1 Terabytes of data per second against them, which is a freaking whole lot, like seriously too much. And the botnet had the capability of sending 1.5. They only got to measure 1.1 but it could have sent 1.5. That's thank you IoT devices.

What happened with WannaCry, some of you remember there was a young researcher in the UK Michael Hudgens, who by pure luck registered a domain that was embedded in the code of the malware, and just by registering the domain name the malware stopped spreading. That is good for him, he didn't know that was going to happen, he became a hero immediately, and then later on people realized that there were things that should be decided by others before he could actually be considered a hero because there were FBI indictments against him. Not indictments, I'm sorry, he was charged in two different times by the FBI right before Defcon of last year and this year again the FBI brought further charges against him. I'm not going to go there. Anyway, it's complicated.

But anyway, he was able to stop the spread of WannaCry just by the creation of that domain. It's unsure on why the operators of WannaCry included that feature in the code. There are some theories but it's not really sure. Probably they were trying to prevent their malware from being analyzed in sandboxes or something, but it's just theories.

Then are the border services. Border services are stress services, are what I've seen in my experience that it's kids who set up basically data attack providers. You can go in provide them with an IP range or a specific IP address, pay whatever fee they are charging and then off it goes. They fire the cannon at the poor victim, but they sell their services. You can test your own network and see if it's properly protected. If it's properly mediating these attacks, but they don't make any sort of validation of whether or not. The test is really on

your own infrastructure, and then if it was there are all these casualties, there's all these side effects you're very likely to take people down while you do that test. So it's not cool, let's say it that way.

I'm just thinking of things that I wish I could share but I don't know if I can. I'm going to put it this way, look for Mirai, infrastructure, domain name, kid and [inaudible] and you'll find -- it's interesting.

These were the malware families that Avalanche was offering to their customers. You could choose any of 20 banking Trojan families. I already mentioned all this, so again go ahead. I mentioned the 64 TLDs that handled 30 000 domains that were taken away of the hands of the criminals, some got suspended, the ones that currently existed at the time of the operation, some got suspended, some got sink-holed.

Sink-holing a domain means that you replace the named server with a named server that's appointed by the court, or that's ordered by law enforcement, or you simply change the named server. Whoever is operating the named server will be able to basically identify the victims. They will be able to tell a lot about the malware itself, and then they will start warranting existence at the time of the operation basically just won't see the light for the next 20 years or so. So the criminals lost complete control of their infrastructure. And there was a lot of things going on. There were a lot of law enforcement agencies, there were a lot of [inaudible] that had to be produced and served as

part of the entire process. It was long and complicated but it was successful in the end.

But it didn't end there, Andromeda was a follow-up to Avalanche. Andromeda was using less domains, I don't know how many, like 3000 or 1500 something like that. But as usual it was the next iteration, the next incarnation if you want, or a child of Avalanche. So also taken care of that was December of 2017. Europol, and the Germans, and the FBI leading all the efforts there again making use of the RSR. I want to stress again how successful the RSR has been. It's a process within ICANN that allows for DDA domains and technically speaking not only DDA domains, to take the domains away from the hands of the criminals.

It's a request of the TLD's that the registries have to submit to ICANN, and basically what it does is that it -- they are asking ICANN for a waiver of the registration fee for those domain names. So ICANN doesn't charge the fee, ICANN waives that fee and they can include in their they can -- Well they basically just take away all their domains that they have bought and that is going to create in the future without having to pay any fee. If it's 800 000 domain names, then that would be a significant fee. So ICANN can say, "Okay this is law enforcement, this is an investigation, this is something that just create them, take them away from the criminals and don't pay me anything because there's no sense in you paying for that," and they just go ahead and it's been quite successful.

Other successful news of the RSR that I can think of are game-overs using crypto locker back in 2014 I think. Game-overs being the most active banking Trojan back then, infrastructure that was operated by a subject who's still the FBI's cyber most wanted. You can look him up, his name is public so I can say it here it's not any secret or anything and I'm not saying anything against him I'm just saying it's the FBI's cyber most wanted. It's them who say it not me. Mikhail Bogachev, that's the subject of interest for the FBI, and he's in Russia, or so people think I don't know.

There's the sided complication with the IoT devices that many are just simply and absolutely insecure listening port 23, Telnet, and 2323. So you can basically compromise them and do with them whatever you want. You can turn them into mail servers, you can turn them into DNS servers, you can do with them anything. So that's an added extra complication. Light bulbs turned into weapons, microwave ovens, refrigerators. It kind of sucks. Seven minutes, thank you Cathy.

Mirai, mostly IoT devices. [Inaudible] use plaintiff's channels, Telnet is just plain text [inaudible]. The administration of those devices is accessible on the [inaudible] admin account and a password. You can find dumps of those default passwords in many places, both open and in not so open places. So even if there was a password set by the hardware producer then it's as if there was nothing because all the passwords are out there and the usernames.

This is going back to WannaCry, when the researcher discovers sink-holes, they erase [inaudible] to control the domain, this is slightly

incorrect and I'm sorry about this. It really wasn't the command control domain. It was this domain that the creators of WannaCry put in there maybe to prevent analysis in sandbox. So when law enforcement or threat researchers have these machines that specifically there for it be compromised with the malwares of their choosing, the certain behaviors would tell the malware to try to delete them itself or what have you.

So it wasn't taking clear command and control and this is when this guy -- Oh well yeah, there were two domains that were -- two or three. It was first one, that most famous kill switch and then a follow-up with one or two extra domains, I don't remember exactly. That maybe was to have the malware detect sandboxes. Okay.

Abusing in the ICANN context. I did mention in the beginning that DNS Abuse is not really well defined, depending on who you ask you'll hear different things. Within the ICANN environment there's always the question of, "What about spam?" That's a really hard question, it makes for fun conversations. There are really opposite views and we're in the middle, so we hope there will be definition soon. There is one thing that the threat research community does observe, and apparently I'm with data to back there their assertions that spam is an indicator of other types of malicious activity. So basically when you see spam there's something else. Spam is a means of delivering phishing for example. Spam is a means of delivering links for malware distribution to get the big things to click somewhere and download the payload, and get compromised. But that conversation isn't going. That's here within the ICANN environment.

Of course if you ask a threat researcher, or law enforcement, for them their definition of DNS Abuse is basically all encompassing. They'll want everything in there, but for us it's narrower than that. So I have to be -- We have to walk a careful line, because when we want to help in my team, in the SSR team, we want to help law enforcement, and we want to help the threat research community, we want IP lawyers to be able to contact the registrars and have things going, but we have to walk that very careful line. Because we can't suggest, or indicate, or we can't let the threat researchers or law enforcement think that certain things are DNS Abuse for us while they're not.

So we have to try to make them understand what can be addressed, and this goes exactly to the point of the ICANN contract. And that's why I mentioned corporate infringement, huh-uh not the path, trademark infringement, huh-uh not the path. And specifically relating to the [inaudible] within the registrar accreditation equipment which is the contract between ICANN and the registrars, don't even think of using it for takedown of pirate websites. That's not the route.

Even for cases when there's phishing but phishing that's not necessarily targeting banking institutions. There's a country where a phishing gang mimics the website of a government entity and they -- what that government entity does, or part of what they do is they sell machinery construction to the public, and by mimicking that website they have been able to trick many citizens of that country to thinking that they are giving a down payment for the [inaudible] bulldozer or what have you, and then the poor buyer doesn't receive any machinery.

That's been a somewhat complicated case to address because the registrar's in a different jurisdiction, they are not familiar with the governments structure within the country where that's happening, they don't know that government agency that does that so they keep going back to the fact that they can't tell who the legitimate -- whether or not it's phishing basically. So it's -- even though it's phishing. But it's understandable that the registrar needs to make sure because there's liability in between and all those normal questions.

So not every case is clear cut, not every case is easy basically. It's easier when it's command and control domains. It's easier when it's DGA domains. The DGA domains since they are simply created by a mathematical function it's just random characters and usually they're not interacting with randomly created domain links. Domain links are created so that we can interact with them. So why would I want to interact with a trainee character long string that was randomly created that I can't even remember? Well maybe there can be someone wanting to use a domain like that, but it's unclear.

With regards to abuse, some topics that are of course of highlight during this ICANN meeting, WHOIS GDPR is of the impact that the implementing -- How the GDPR impacts the availability of WHOIS information that's a technical conversation that law enforcement and the threat researches are having with the ICANN community. Today there's participation of some high level representatives from Interpol and Europol during the HLGM, the High Level Governmental Meeting that's going to be interesting if you're interested you can follow that.

Well, in the schedule you'll find all the WHOIS EDPR related sessions of which I think there was quite a few.

RDAP implementations from our perspective, we need to have the threat researchers and the law enforcement agents become familiar with RDAP before it gets implemented, because they need to know how it's going to look like, they need to understand it, they need to be familiar with it. Because one day the WHOIS particle is going to go away and it's only going to be RDAP, and they have to know by then how to use it or else they'll just going to be delayed in their investigations and they're going to be interrupted for a while until they learn. So they need to start learning now, and if anyone of you is here a researcher -- threat researcher or law enforcement by all means please, please, please read and play with RDAP. Verisign has an autopilot, some others may as well so just keep an eye on that.

And then idea and implementations. From an outsiders perspective I included idea and implementations because there is that -- maybe it's not a trend it's just an observation by some in the security community and I know we're out of time. Yeah okay, last thing. Domains, idea and domains using more than one string are being seen in abuse. Some blog posts have been posted particularly by Farsight Security policies company, they are looking at that this and the findings are sadly very cool.

Sadly because it's criminals being clever and registering domains that include, let's say, Latin and Cyrillic, or Greek and Hebrew, and just making samples up to make them look like the real brands and then

phishing the victims. So that's something for the anti-abuse community to keep an eye on. Just see what's going on where the guidelines go to and just for awareness I guess.

I already talked about the Hyderabad GAC Communiqué, so we can continue. PSWG, please remember the PSWG is basically the Public Safety Working Group that's if you want the law enforcement space within ICANN, they exist within the GAC. These are some of their topics of interest, GDPR of course, WHOIS accuracy, whatever that means now under the GDPR, that's a conversation in and of itself. I'm not going to go there. Contractual compliance and the DNS Abuse at large. I already talked about this. Then some provisions in the registrar increment that have to do with anti-abuse. You can check this up, this is published on their public schedule so you can go and look it up.

Again, these sessions are published in the schedule so you can look them up. And that's it. Thank you so much, I hope it was of interest. If anyone has questions since I think we're out of time, then I'll just be outside or in the back. Thank you so much.

[END OF TRANSCRIPTION]