COPENHAGEN – How It Works: Understanding DNS Abuse
Monday, March 13, 2017 – 13:45 to 15:00 CET
ICANN58 | Copenhagen, Denmark

UNIDENTIFIED MALE: Please feel free to come up to the tables. We have power and microphones up here. This is a pretty informal session and we encourage dialogue, so if you want to be at the table and have power, come up and join us.

STEVE CONTE: We're just about to get started but again we've got plenty of table space up here. The table has some microphones if you wish to engage in the dialogue and most importantly power, please come up to the table. It's going to be very informal. We're hoping to have a dialogue and not just point you at the side deck, so come join us.

All right, with that, we're going to get started. Thank you. I hope you all had a good lunch. For the past couple of days, I've been thinking, "What a health conscious country this is. Everyone is drinking milk." But it's water. I'm sure it's a health conscious country.

So, this is the second of today's series for the How It Works tutorials. It says John Crain but that's not John Crain. We have

Carlos Alvarez here from our SS, Security, Stability and Resiliency Team, and he is going to produce the presentation today. If you have any questions at all, if you're not at a microphone, please raise your hand. Carlos, you have the only handheld mic, so you get to go play [inaudible] walk out there. If you are at the table, please raise your hand. Carlos will acknowledge you and we can have the dialogue that way.

We do have remote participation. So, if you do have a question, please don't shout it out. We're trying to capture it on Adobe Connect as well because we'd like our remote users to be able to engage and listen to the conversation, too.

So with that, I'm going to pass it over to Carlos.

CARLOS ALVAREZ: Thanks, Steve.

Hello, everyone, and thank you for being in here. This is going to be for the most part, an introductory level session. We're going to talk about DNS abuse, what it is, what it is not, whether it is defined. And we're going to talk about a little bit about what ICANN does with regards to DNS abuse very, very briefly.

As Steve said, I'm not John Crain. I work with John Crain. I work for him. I'm a senior manager with the Security, Stability and

Resiliency Team and I live in Miami, so Copenhagen is cold, yeah, but anyway.

The idea is to have you all interested in this topic. There are many constituencies and interest groups within the ICANN community that have a high level of interest in DNS abuse. We're going to see how it looks like, what it means, what it doesn't mean and how it can be addressed. So, let's get started.

What is DNS abuse? What it is within the ICANN context? We're going to talk a little bit about examples of DNS abuse and then there's going to be some links that you can go to afterwards to learn a bit more about abuse. The slides are going to be made available in the schedule.

UNIDENTIFIED MALE: [Inaudible].

CARLOS ALVAREZ: Oh, they're up already, so you can go in afterwards. For now, just pay attention to me and just read a little more. But hopefully you'll be interested and continue exploring on these topics because they're really interesting.

So, what is DNS abuse? There's been a lot of discussions for sometime already on what it is, on what it should look like.

There's no clear agreement within the ICANN Community of what DNS abuse actually is. Some parts of the ICANN Community would like DNS abuse to have a very broad scope, some other parts would prefer that DNS abuse be very narrow. We're just expecting the community to inform ICANN of what DNS abuse should be defined or if not, a proper definition at least define what DNS abuse is, what types of activity can be considered as DNS abuse.

DNS abuse covers a wide range of activities. We'll see those in a minute. There's no globally accepted definition but there are definitional variants. Some of which can be considered heretical or heretic by some here in this community such as general cyber crime and hacking, which might make some itch around these tables. That's not what we're talking about here, not generally a cyber crime and not general hacking, too. That's not really what we're talking about here. It's much more specific than that.

In the past, ICANN has talked about malicious conduct in a study conducted by the University of Rome and the Global Cyber Security Center have categorized DNS abuse within three precisely categories: data corruption, denial of service and privacy.

I'm going to read this. DNS abuse refers to intentionally deceptive, conniving or conniving or unsolicited activities that

actively make use of the DNS and/or the procedures used to register or resolve domain names. And this goes to what we care about.

As ICANN, we care about the provision of the domain registration services and we care about the resolution of domain names. We care about the – there are solution process of domain names and the infrastructure associated with the resolution of domain names. So, this is more aligned with what we really talk about here.

In simple terms or simpler terms, DNS abuse refers to anything that either directly abuses the DNS infrastructure like those large DDoS attacks that make use of the DNS as a vector to launch massive amounts of data against a target. Or misuses the DNS protocol and the registration domain process for malicious processes, like when criminals register large amounts of domain names to provide command and control for their botnets – and we're going to talk about botnets in a bit.

Maliciously registered domain names – sometimes, criminals do actually go on registered domain names, from the very beginning with a malicious intent with the intent to harm someone and commit some kind of malicious activity. If you notice, I'm not using the word illegal activity or illegal better, I'm talking about malicious because we don't want to get into the

discussion of whether it's illegal in one jurisdiction and not in another jurisdiction. It's just generally speaking malicious activity or abusive activity that can harm or that does actually harm users.

These are some examples here. Phishing – we'll talk about phishing in a little bit. Malware – botnet command and control. Data exfiltration. This is a somewhat novel way for criminals to use DNS infrastructure to actually exfiltrate data from their target victims, networks. Malware distribution, exploit attacks, scams, counterfeit goods, illegal pharmacy and infrastructure like when the criminals compromise ISP name resolution servers to victimize their ISP's customers.

Abuses for other people's domains and DNS – this is more related to when criminals compromise other people's domain names or servers. And, here it's a little more complicated than when it's an obvious malicious registration because the registrant registered the domain with a licit and legitimate purpose and using the domain name with that licit and legitimate purpose. However, the criminal gain control of that domain or of that server so it puts an extra burden on the people that are the doing the investigation to address or mediate or contain the threat because they can't just go after a suspension. They can just go after a takedown because it might cause third

party harm and that's something that I need to take into account and prevent for the most part.

Hosting criminal DNS infrastructure, domain name hijacking, hacktivism, tunneling, host file modification – we're going to talk about all these in a little bit. It's jut listing them right here. Changing default resolvers, there is an interesting case that was addressed by law enforcement. I think it was 2012, several law enforcement agencies from different countries conducted a big operation towards a large botnet that was called DNSChanger. The malware type was called DNSChanger – we'll talk about it in a little bit.

What it did basically just going a bit ahead of myself was it changed the DNS configuration within the compromised devices, so that the DNS queries wouldn't be sent to the legitimate DNS servers but rather to name resolution servers operated by the criminals. That was mean and that was evil, and a lot of people got harmed by those criminals.

Poisoning resolver or ISP servers – as I mentioned, when the criminals get a hold of the ISP name resolution servers and modify the resolution formation that's within their cache memory, which allows the criminals to direct people's traffic to the servers that are again operated by the criminals. This allows for them to steal access credentials to banking websites to clone

whatever news website you want to look at and have them implant malware, and have that malware injected into your machine. So, a number of bad things can happen via poison resolver/ISP cache memory.

These are some descriptions of how the bad guys can actually attack the DNS.

Basic cache poisoning – it can happen at several levels, it could happen at the device level or at the ISP name resolution server, which are their recursive servers. Basically, what it is what I just mentioned 30 seconds ago, the criminals compromise the user's device or the ISP server and change the cache memory, and modify the IP addresses.

If you guys remember how resolution works, each device has a file called the host file that includes their associations between domain names and IP addresses, and a similar memory is within the resolution servers of the recursive servers that points each domain name to the intended IP address. Well, what the criminals do is change those IP addresses and include IP addresses for servers that they own or operate. The result of that is, as I was mentioning that the user's traffic gets directed to infrastructure paraded by the criminals, which is simply criminal infrastructure. That's there to inflict some form of harm either by stealing banking credentials or injecting malware or simply

sneaking to people's traffic and just conduct an espionage whatever you would like to do with that.

Indirection attack – DNSChanger, this is a little bit different. This is the example that I was mentioning. The criminals were able to modify the actual DNS configuration within the compromised devices. And what that meant was that – let's say my laptop, I have configured my laptop so that it sends DNS queries to Google say 8.8.8.8. The criminals would change that configuration so that DNS queries coming from my laptop would be sent to 1.1.1.1, which will be their own name resolution server. They would have configured that server so that responses coming to my device, to my laptop would be pointing towards servers that they operate.

So, if I want to visit any news website for example let's say CNN.com, as just any example, my browser would resolve, meaning that it would take me to visit www.CNN.com. It will be an almost real-time image of their real website. What I would see in the browser in the URL would be CNN.com but it would be hosted at the criminal's server.

And so, the criminals would own that server. They could do whatever they want with that content. They could replace ads and have their own ads there, so that they would generate income for them. They could also include malware in the ads or

in the graphics or in downloadable PDFs. They could also redirect banking website's traffic.

So, if I wanted to visit BankofAmerica.com, I would go to a website that would – in the URL it would actually be www.BankofAmerica.com and I would see an exact copy of BankofAmerica.com but I wouldn't be sending my username and password to Bank of America. I would be sending my credentials to the criminal server. As a result of course, five minutes afterwards, my savings account would be wiped.

So, the FBI and other agencies took action against this infrastructure. As I said, I think it was 2012 that is a large operation that included people in many countries. They got a hold of their servers that were serving as command and control for this botnet.

My throat is not behaving well, so apologies. I may have to put some mint in my mouth. I have them ready right there just in case. But if I cough, so just bear with me.

Then, law enforcement was able to get a hold of those servers. Those servers were giving administration to the Internet Systems Consortium, which was appointed by the court to continue the operation of those servers. And the reason why those servers have to continue to be operated was that because there were hundreds of thousands of compromised machines in

many, many countries. If they were turned off, the users that were using those compromised devices would think that they had lost access to the Internet because there would have been no resolution.

In other words, if a user using one of these compromised devices wanted to go to CNN.com, it would send the DNS query to a server that was shut down. There would be no response, so the browser wouldn't take the user to CNN.com. As a result, the user would think that he had lost Internet connectivity and that wasn't the idea.

So, those servers continued to be operated by the good side of the force, the light side of the force better. They identified all the compromised IP addresses or the IP addresses from which calls to the command and control servers were being made. They segregated per country when you have the IP address as you can now, which ISP for the most part, which ISP is actually has received that IP address within it's allocated blocks or ranges, and you can now reach country that ISP is located in, at least that's an idea. And that through the national CERTs try to have those, the CERTs are the Computer Emergency Response Teams have those CERTs approach the ISPs and ask them to ask the customers to clean their machines.

So, it is a long chain. There was a lot of awareness, a lot of education involved and at certain point, those servers were shutdown with enough warning given to the community, to the Internet community in general that those servers were being shutdown.

Distributed Denial of Service attacks – a resource depletion attacks where thousands of bots send DNS queries to a target name server. I think that first [beak] attack of this type happened in 2013 against Spamhaus. Spamhaus is – it's up at the security group I would call them. They do a lot of work, of course, an anti-spam but also they do investigations in malware, phishing and other kinds of malicious activity. And, they oust people when they identify someone that they think is conducting a malicious campaign that's very harmful, they publish their information on Spamhaus's website.

So they did in an instance back then. And the person that got ousted by Spamhaus got quite upset. This person which was a very young man had his group of friends directing attack against Spamhaus. Back then, it reached I think 330 gigabyte per second, which is quite a lot. And, it took Spamhaus's name servers down for sometime. They had to jump on CloudFlare and they had to do a lot of mitigation. But the attack against Spamhaus that only used the DNS for resource depletion was successful for sometime.

The result of that attack, the way it works and again going a bit ahead of myself is that the criminals are operating a large botnet. A botnet, for those of you who may not know is a network of compromised devices that answer the orders given by the person that's controlling the malware. That person that controls the malware is generally known as the botmaster and that person can inject more code in the compromised devices, send commands to the devices and have them attack someone else, exfiltrate information from those devices. So, they can do many, many things. They had the malware already implanted in through their command and control servers.

So, in this case, what the attacker [see] was that they ordered the compromised machines that were thousands spread throughout many countries. They ordered those machines to send DNS queries to servers that are known as open resolvers. Those open resolvers are recursive resolvers that are operated by ISPs in many countries. They have a particularity and that is that they answered DNS queries from anyone in the world.

In the ideal scenario, a recursive server or an ISP name server resolution server would only answer queries sent from the ISPs [on] customers. So if I run an ISP and I know that my customers receive from me, IP address is only from this specific range, in an ideal, super, super ideal theoretical scenario, my resolution

server would not answer queries sent from customers outside of those IP address ranges.

There's another scenario that's fair and that's 100% okay. It's the one that I mentioned, like Google say 8.8.8.8, they answer queries from anyone but they are very well managed, so they won't allow these sorts of attacks to take place through their infrastructure.

So those open resolvers were used by the attackers, they compromise machines, ended up sending thousands of DNS queries to many thousands of open resolvers and all those queries were indicating the servers that each query had been sent by Spamhaus's name server. You can forge the DNS queries. You can forge any IP packet that's traveling through a router and spoof the origin of the source IP address.

So by spoofing the source IP address, the attackers got all the open resolvers or the name resolution servers to send the responses which were much larger than the queries against Spamhaus's name server and it crashed. Boof. And the result of that was that people trying to look up the IP address for Spamhaus's resources, their website, FTP, etc. could not find anything. So, that is the first large, large case where this sort of attack took place.

Well, a similar situation here, DDoS amplification – thousands of bots issue queries that evoke a very large response message. They all spoof the address of a targeted name server and the targeted name server is flooded with a very large DNS response requested by the compromised computers. So, it's very, very similar here.

The exploitation attacks – they exploit a software flaw that causes DNS servers software to fail or behave in an unintended way. [Box] in the system basically.

Abuse in an ICANN context – there have been for sometime already discussions within the ICANN Community with regards to what DNS abuse should actually be. There are WHOIS accuracy discussions related to that. Those discussions seemed to be endless. They have been there I think since the beginning of time. Always there are of course public safety issues always associated DNS abuse.

DNS abuse within a kind of strict kind of way can be understood basically as – and this is for the GAC, the Government Advisory Committee that issued – I think it was the Beijing communiqué, April of 2013, I think that was Beijing. I may be wrong. A communiqué indicating that abuse in the ICANN context should be or they understood it as being phishing, farming, malware distribution and botnet command and control.

And then of course malicious registrations that are meant to facilitate child abuse materials [inaudible] and distribution would fall in that description and this is just my very personal [inaudible] but many people agree.

Yeah, Beijing – what the GAC said back there was that registry operators would ensure that terms of use for registrants include prohibitions against the distribution of malware, the operation of botnets, phishing, piracy, trademark or copyright infringement, fraudulent or deceptive practices, counterfeiting or otherwise engaging in activity contrary to applicable law.

The input is a lot. This is not set in stone. There are discussions about this. There are people who don't want piracy, trademark or copyright infringement there. This question about what is applicable law is always there. It might be illegal in one country but it may not be illegal in the neighboring countries, so there are all these questions about that.

And, please if any, if you have any questions, please feel free. The idea is to make this more of a dialogue rather than a monologue. I hope I'm entertaining enough that you guys won't fall asleep right after lunch. I have an energy drink here but it's just one so it's not going to be enough for everyone. So, yeah.

Something interesting that has to do directly with DNS abuse or how it's handled or should be handled or is to be handled by

ICANN is the creation by the law enforcement community of what is known as the Public Safety Working Group. The law enforcement community had already participated within ICANN for a very long time for many years but they didn't have a formal spot where to sit. So they were there but no one was really sure where that was and they were providing advice and they were providing guidance to ICANN. And I'm using the word guidance in the most informal way possible to the ICANN Community in general and I'm not just talking about ICANN, the organization.

And something that they had been working on as of recently and maybe – well, I mentioned they GAC communiqué of Beijing in 2013 but more recently than that, probably last year at some point, they started actively talking, discussing about what DNS abuse is. And that is caution of whether it should be defined or if just it would be better to simply identify the activities that might correspond to DNS abuse without trying to define it.

Because having people agree on that definition is going to be hard like seriously hard. And the idea is to address known forms of abusive activity that make use or rather abuse registration services and abuse resolution services. The idea is to address those types of activity, and not just seeing the discussion for years and years.

Some contractual elements that are there that were included. This provision here was included in the new gTLD Registry Agreement. That's the base registry agreement that all registries nowadays have to sign. It's Specification 6 that provides that registry operators must appoint and publish an abuse point of contact. And that talks about malicious use of orphan glue records. Steve, if you want to clarify or if you would like to add something a little bit about malicious use of orphan glue records.

STEVE CONTE: You kind of put me in the spot now. I don't have anything else that – what John has been involved with more than I have, so I hope [inaudible].

CARLOS ALVAREZ: Okay. Then Spec 11.3 – there's been a lot of discussions about Specification 11.3b. There are two things in this Specification 11 to the Registry Agreement. There's one part that talks about the security framework that it's aimed at indicating how registries should or could act on malicious domains that are within their zones, that means, every domain that's within .com, every domain that's within .guru and so on.

And Spec 11.3b that provides that registry operators must analyze their domains within their zone to identify malicious registrations and reporting them to ICANN. That's what it refers to. Identify them and report statistics to ICANN. So, that's there and this is part of the public interest commitments.

There are some registries that have stronger PICs as they are called, the Public Interest Commitments in their agreements. And then, according to their PICs they have to or would have to take more provocative fashion towards addressing abuse but that's not for all the registries. It's just those that included that like extended version of Spec 11.3 and Spec 11.4 I guess.

Then the RAA, that's the Registrar Accreditation Agreement, which is the agreement between ICANN and the registrars. Going back a little bit in time, the first version of the agreement between ICANN and the registrars was from 2001 where those that negotiated the agreement just had no way of foreseeing what was going to be of the DNS, of the name space in general with regards to legitimate uses and malicious uses of their domain names.

That agreement included no provision whatsoever related to abuse, like none, zero, zero, [inaudible], nothing. Then, the following version 2009 was a little bit better. It included some more provisions that allowed ICANN to take more action but still

didn't refer to abuse specifically. And, some sectors in the community felt that a stronger agreement, some stronger provisions were needed there in the – I think it was in Costa Rica 2012 or Senegal also 2012.

The law enforcement community through the GAC, through the Government Advisory Committee came up with 12 recommendations that were provided to the Board of Directors. And those 12 recommendations, the Board instructed the staff, the ICANN organization, to initiate discussions, negotiations better with the registrars towards implementing those 12 recommendations in the RAA. Those negotiations took place for a long time.

And, of course, both sides have had to make compromises, both sides had to give up some of what they were looking for because it is a negotiation. ICANN as you may know does not impose the agreements on the registries or the registrars. These are private contracts that are negotiated with the other parties. So we have to actually sit with them and negotiate. It's not imposed. ICANN is not a regulator.

So then I think one of the most – not the most but one of the most relevant provisions in anti-abuse that ended up being included in the RAA was Section 3.18 that on one hand provides that the registrars have to publish an abuse point of contact,

which didn't exist. The importance of an abuse point of contact is that on one hand, the security community, which is very active in identifying malicious domain names or domain names being used with malicious purposes need to know where to go. If they go to registrar's finance person or if they go to a customer service representative, they might not get a proper answer because that person has no idea what the security researchers are talking about. So, the appointment of the abuse point of contact is intended to provide a proper channel for the reports of abuse to submitted to the registrars.

Then, it also provides that registrars have a duty to investigate reports of abuse. The RAA provides that registrars shall take reasonable and prompt steps to investigate and respond appropriately to any reports of abuse. This of course is lawyer drafting, reasonable and prompt steps. If you ask an engineer what that means, they will just like, "Pfft. We don't like that." But the lawyers, on the other hand, say that's the only way to draft it.

So, there's that tension between the security community and the negotiators that came up with this [fording]. And the reason is because sometimes what to a security researcher should be reasonable is not to the registrar that receive the complaint or the report of abuse and not [inaudible]. So, there is that tension. And, it's very hard to go about that tension just because of the

way contracts are drafted. It couldn't go down to details on a case-by-case basis. That's not the intent and that's not what a contract should necessarily look like. So, it's a point of tension I think.

There's also suggestions that we – we call our team the Security, Stability and Resiliency team, we call ourselves the SSR team and we do a lot of engaging and liaising between the operational security community on the registrars and the law enforcement community as well. We always try to have the security researchers understand that they have to do their part as well, that they have to provide enough information up front to the registrars when they submit reports of abuse.

Sometimes they get frustrated when they submit reports of abuse. But when you go and look at that report, how did you even expect the registrar to act on it? There's by no chance enough information there for the registrar to figure out what you are talking about. It happens occasionally.

There are some in the security community that understand this very well that send the reports in a very complete fashion, sometimes with too much information, which I guess is not that bad.

So, there's that. I guess what I'm trying to say is that there's a lot of work being done in this regard to have both sides, the

operation security community and the registrars both sides be better at this. Of course, there are registrars that are really good at this. Not that they're not good but it's an ongoing process. Everyone is learning on this I guess.

Interestingly, it's not mentioned here but also the RAA provides that this point of contact for the registrars must respond within 24 hours when they receive reports of abuse submitted by law enforcement from their own jurisdiction and the response not only has to be within that timeframe but it also has to come from someone with authority to act on the matter. So, it's not just any tier 1 person that's handling the reports of abuse. It has to be someone who can actually suspend a domain if it's needed.

This person needs understandable why it's there. However, in practice it makes things a little bit complicated when its law enforcement from one country who are saying malicious registrations in a registrar that's in a different jurisdiction because they have to reach out and obtain cooperation from law enforcement in that other jurisdiction. So the mechanics of it may not always be helpful but it is what it is. It's what's out there. It's what's in the RAA and it's what we have to go by. So, it just is what it is.

There are some sessions this week that might be of interest to you. This was today. This is happening right now. So, because you're here, it means that – the recording is going to be posted at some point. I'm not sure when. It usually takes one or two weeks or so, so let me go back and see that recording. This is going to be a very interesting discussion. Many sites of the ICANN Org and the ICANN Community are going to be talking about Mitigation: Prevention and Response towards DNS Abuse.

Then tomorrow at 11 session from ICANN's Global Domains Division about the Statistical Analysis of DNS Abuse in gTLDs. That CCTRT, that's the Competition, Consumer Trust & Consumer Choice Review Team commissioned a study on DNS abuse within the new gTLD space.

The third party that was hired to produce that study is going to give the community a preview, so it's interesting. Yes, sir?

UNIDENTIFIED MALE: It seems on the talk that the answer to DNS abuse is procedural movement but what about the technological aspects, like DNSSEC for [inaudible] is solving some of the things that you mentioned, so what are the technological tools or means that can be used?

CARLOS ALVAREZ: There are different types of DNS abuse under different types of technological responses to those types of abuse. DNSSEC is one of them. Some of you might remember what DNSSEC is. In a very, very simple form, it prevents spoofing. Make sure that when you go to a website, it is the website that you intended to go to, in a very simple way to put it.

But there are way much more technological responses to threats using the DNS. There's SPF, DKIM, DMARC and a number of others. But the thing is that it's up to the actual community to implement those solutions. There are sometimes in which those solutions don't really have to do with the ICANN mandate.

For example, BCP38. BCP38 is what's called the Best Current Practice that was defined as such by IETF I think in '98. And it addresses what I was mentioning with regards to the lack of source address validation. It provides good ways to go about solving that issue. It indicates how ISPs can filter out packets coming with spoof IP addresses but there's little that ICANN can do about that about the fact that ISPs don't actually implement BCP38. It's not within ICANN's remit. So, there's the solution, there's the problem, there's the solution, it's voluntary implementation and it's outside of our remit. So, we can't go and act on that.

Speaking of DNSSEC, registrars and registries – the new gTLD Registry Agreement on registrars in the 2013 Registrar Agreement that they have with ICANN, they have the obligation to implement DNSSEC. And it's taken seriously. I just had a conversation with a registrar whose name I'm not going to disclose. They are still under the 2009 RAA and they are willing to be able to sign the 2013 version but they are fearful because they need to implement DNSSEC. And they need resources for that so they are asking ICANN to come up with a plan for them to implement DNSSEC in a few months while still being able to sign the new version of the RAA soon so that their accreditation won't expire.

The PSWG, that's the Public Safety Working Group. It's going to take place tomorrow as well at 1:45. It's going to be a presentation to the GAC Plenary. This is interesting because it's actually the Public Safety people that are going to present to the GAC what they are working on, what they are saying, what they're focusing, how they are seeing the work of the CCRT, the Competition, Consumer Trust & Consumer Choice.

And, there's that somewhat of a focus on the New gTLD Program and anti-abuse in general. So, for those of you interested in abuse, this will be a very interesting session of course.

Now, talking about some examples of how DNS abuse can look like.

Malware comes from malicious software. It's basically a software used to compromise devices. I don't think that needs much explanation. However, nowadays and for sometime already, malware is used by the criminals so that when the devices get compromised, the first thing that device does is make a DNS query. That's the choice for bad actors to control their botnets via the DNS. So, there's a DNS component in most every form of malware nowadays. As soon as the device gets compromised, it makes a DNS abuse query. That's the norm.

Basically, the reason why is because the DNS traffic goes through a port that's port 53 and it can be redirected in the internal network, meaning that the engineer, the system, the systems administrator in the company can redirect it so that inside their boundaries, the traffic goes through a different port. However, not all engineers are up to that.

Of course, it just can't be blocked. If they block port 53 without doing their redirection, then users will think that they have no Internet connectivity and they won't like that.

So, port 53 is almost mostly certainly still being used by most everyone for DNS and it won't be blocked. And for the most part, it's usually not monitored or analyzed, which means that even if

the company has a sound security infrastructure, it's likely that they want to be looking at the DNS traffic, so they will miss. They will still get compromised. The anti-malware products and the firewalls, and all the stuff that one can implement in one's own network will not be enough. It's just a matter of when it's going to happen nor if it's going to happen. And when they get compromised, they won't realized that they were compromised unless they have good monitoring and analysis of DNS traffic associated with rules that trigger alerts and some more things that they could do, only out of monitoring DNS traffic.

Botnet Command and Control – there you go. It had to happen. So, botnet comes from roBOT NETwork. As I was mentioning, it's just a very – well, it can be large or small but they usually are pretty large network of compromised devices that can be controlled remotely by the attackers.

There was one instance that was particularly evil where the criminals after the machines got compromised with a form of malware whose name I just forgot – oh man, that's bad. Well, I'm sorry that I forgot the name of that form of malware. They injected a second form of malware that was called Game over Zeus. It's a Trojan banking. If I'm not wrong, it is the most active Trojan banking in 2014 or 2015.

And after they, injected a third form of malware that was CryptoLocker, which is the great, great, great, great, great grandfather of the Ransomeware types that are out there today.

So, with the first form of malware they implanted a software that allowed them to launch attacks against third parties. With the second form of malware, they stole access credentials to banking and financial services of their victims and wiped out their accounts. And with the third form of malware, they encrypted the important files of the users and then asked for a ransom to be paid in the form of bitcoins. So the producers were now in a really, really bad situation.

All this happened as a result of the criminals being able to use the port 53 as a communications channel using the DNS. Through the DNS, the criminals receive the 80 phone home calls from their compromised devices to their command and control servers, send responses, inject more malware.

It's actually very cool to see when you are capturing data, when you're capturing traffic and you isolate the DNS traffic. It's very cool to see the interaction between the compromised devices and the command and control servers. You see everything that's going on there. You see the commands coming in that the compromised device starting to scam the neighboring machines in the internal network. You see everything that goes between

the command and control server and the compromised machine.

Why are botnets interesting? Why are they not? Since late 2008 I think, one of the most pervasive forms of malware was detected by the entire security community. It was called Conficker and there are some variance of Conficker that are still out there today that great, great, great, great, great grandsons of this bad Conficker back then.

What the criminals were doing with the Conficker was that they had scripts in the malware that would have the botnet automatically create strings that were then registered as domain names. That's what's called a domain generation algorithm. So, the botnet didn't need human intervention to have a pool of available domains everyday for command and control of the entire criminal infrastructure. They had associated stolen credit card information and many resources so that they could just run the process automatically without needing to have someone there [sitting] and registering the domain names.

Conficker was pretty evil because it was registering hundreds and thousands as John had it here. He mentions more than 50,000 domains per day across more than 100 TLDs. So the response was pretty complicated because it's one thing if you are talking to a one registry operator about taking down

malicious infrastructure. It's very different if you have to speak with 100 TLD operators, most of which were ccTLDs with their own policies, with their own sovereignty that ICANN couldn't step over at all, of course, ICANN has to respect the ccTLD sovereignty. So, it is to say very complex scenario. It resolved but it took time and a lot of discussions.

These are some examples of domains that were used by Conficker in 2013. Just random characters as randomized as the bad guys have them be according to how they write their scripts. They can be 32, 64 – no, not 64, however long they want them, the domains and include only alpha characters, numeric characters makes whatever they want.

So, of course, if a company or a government entity is actually monitoring and analyzing DNS traffic and they have good [inaudible] in there and they suddenly see machines calling out to domain names that are nothing but randomly selected alphanumeric characters, the probability that that's going to be a DGA domain, a algorithmically generated domain name by a botnet is going to be very high. So, blocking it might save the company a big headache. Blocking, not allowing the DNS query to go out so that the command and control server doesn't know that that machine is compromised.

When a botnet is identified, one of the ideas is to take control away from the bad guys so that they lose their infrastructure. Aside of course from law enforcement interest and identifying and capturing, and prosecuting the criminals from an infrastructure perspective, the idea is to take it away from the bad guys. In terms of the protecting the system, the DNS as a whole and disability of the Internet, etc. etc.

And to achieve that, what has to be done is to take all the DGA domains away from the criminals, all the domains that are going to be used at one point or another by the criminals for command and control of the botnet.

There was an operation that was led by the German police and Europol in December of last year. It went against the platform, a criminal platform that was called Avalanche. In that operation, 800,000 domains were taken away from the bad guys and the way it works is that the security researchers deciphered the domain generation algorithm and by deciphering it, they are able to tell all of the domains that are going to be registered by that botnet for the next 10 years, 20 years. It's just a mathematical functions. You can just let it run. You'll get thousands of domains.

The way it works is that ICANN as a result of the Conficker threat, their community came up with – it's not a process of such. We

referred to it informally as a process but it's really called simply an Emergency Registry Security Request or ERSR where their registries ask ICANN to wave their compliance with certain obligations in their registry agreement.

Through the ERSR, the TLDs that were going to have domains registered within this 800, they asked ICANN for a waiver and ICANN allows them on the registries. You can either say that they preregistered those domains to themselves or just blocked them from being registered so that when the botnet attempts the registration, it won't happen. And as a result, the criminals lose control of their infrastructure.

Yes, sir?

UNIDENTIFIED MALE: How much time it will take to block the 800,000 domain names?

CARLOS ALVAREZ: The preparation for an operation like this takes a lot of time. If you include the analysis of the malware samples and then deciphering the actual DGA script, it takes a lot of time. Then, the way it usually works is that the law enforcement obtains court orders that are addressed to their registry operators. And, running that, the actual ERSR takes a couple days, not more than that. But for the registries to include those lists of name in

their block list or just preemptively [inaudible] them, however you want to put it, it's just a matter of minutes. Yeah.

Of course, the bad guys only need to maintain control of their botnet. They only need one successful registration to be available per day. So, the analysis has to be complete and no one domain can escape because just one can allow the criminals to continue using their infrastructure.

Then of course, there is the problem that comes when criminals will like what they do and they make a lot of money, and they make enough money to come up with newer versions of their malware and their botnets. So, like what happened, what happened with him? This is public, so I'm not making anything up. This person is wanted by the FBI who was according to the FBI, he's the suspect. He's suspect of having operated the Game over Zeus and encrypt to local infrastructures. And as long as all the infrastructures was taken away from him, it only took him a few months to come up with further versions of the malware. So, he kept on going with his business [inaudible].

But at least, many users were protected from harm. As people in the security community put it, part of the idea is to keep the bad guys a really bad time, give them a headache, like one hell of a headache. And that was achieved which is not the only thing but he had to restart his evil business to putting it away.

We're going to talk a little bit about DDoS attacks, Distributed Denial of Service attacks. Now, that there's the word "distributed" because of that, the large amount of compromised devices that can be part of the botnet and their geographical dispersion coupled that with all the amount of open resolvers that are out there that are operated by ISPs that may not even know that they should be more, be better at managing their name resolution servers. Maybe they just don't even know. It can happen many, many ISPs out there are mom and pop shops. Many are sound companies with large corporate structures but many are not. It's of course hard to tell.

Because of the way the protocols work, you can send a very small query that would be just one line of characters, just clear text, just one line and have a response. That would be 2.7 megabytes, which is more than a thousand times the size of the query.

So, as I was saying, when you have hundreds of thousands of compromised devices making DNS queries to thousands of DNS open resolvers that are going to send thousands and thousands of responses, each 2.7 or more megabytes to one specific device, then the result of course that the poor target is going to be shut down pretty quickly.

This is how it looks like. This is the compromised. This is the attacker asking the compromised machines or the zombies to send their DNS abuse queries to the open resolvers, the open recursive servers. What these queries are telling the recursive servers is that those queries are not being sent by them but by the victim. So these queries indicate that the responses are being sent by 10.10.1.1, which is this guy here. So these servers will unavoidably send all the responses to this guy and take him down. It's off.

And remember, it's hundreds of thousands of zombies that can be out there in any specific botnet and there are – I don't know the actual date but it's somewhere around 30,000 open resolvers out there. There's a guy in the security community, Jared Maunch who runs the Open Resolver Project. He has published his data and he continues to run that project on a continuous basis. What he does basically is that he scans the entire IPv4 space looking for DNS servers and sends them queries to see if they are going to respond. That's an open resolver with – and he keeps track of that, of those responses. It's pretty cool. It's interesting.

This is how DigitalAttackMap.com portrays large DDoS attacks. It basically shows victim location servers that are spitting out a lot of bad traffic to victims. So it's just a graphic representation of how a DDoS attack might look like. There are other maps like

this or hit maps if you want. They just call it a visual way to show how one of these attacks might look like.

Phishing and Spear-Phishing. Phishing – I don't think it's very needed to explain too much about it. It's basically luring a user to do what the attacker wants. Phishing can be understood in a narrow way of saying that it's a criminal that sends an e-mail claiming to come from a bank but it can also be understand in a broader way as a criminal sending e-mail to potential victims in the hope that they will be lured in clicking or giving up their information or something. And from then on whatever the criminal wants to achieve implant malware or extract their credentials or just – who knows. That's up to the criminal

Spear-Phishing, which is targeted phishing. Spear-Phishing is actually that probably the most concerning type of threat because it's the perfect way for criminals to gain access to your networks. Say if someone wants to compromise Göran, our CEO's device from him, move on to the internal network. Nowadays, they want to send him an e-mail. What they'll do is that they will conduct social engineering campaign if you want against us, against the ICANN Organization, try to figure out who works directly with him, their patterns, who they talk to, what they're interested on, at what time they come into the office, who usually writes to them at what time and so forth. And, send them, one of his direct collaborators an e-mail claiming to come

from some contractual or what have you. And ask them, "Dear Steve, I am aware that your CEO is expecting this very urgent information. I need you to please pass it onto him because it's really urgent. We are about to sign the contract but we need him to review these and give his okay."

Steve will, "Oh, my God. Yeah, I have to." And sends it to Göran. And pfft. Göran will see the e-mail coming from Steve, he'll open the PDF and we'll all be dead basically.

That's happening in many organizations. It goes unnoticed and the criminals are just using domains to conduct these sorts of campaigns. So, heads up.

Actually, the Corporate Cyber Security Team at ICANN is always conducting awareness campaigns on us on the staff and they're all really good. Yeah, they sometimes fool us because they're really good at making us fall. And of course, the premise that we come from is that everyone will fall no matter who it is. It can be the CTO. It can be a security guy. It doesn't matter.

If you're tired, if it's a Friday night, if you just want to go to bed and you get an e-mail saying that something is urgent and it needs to be reviewed and you just want to go to bed because you're really tired, you're going to click. So, yeah.

Well, when I see this, it made me kind of laugh and felt a little sad because this is the government of our department or a state in Columbia. I am originally from Columbia. So, evidently, these guys have a compromised server that some criminals are using to send phish, which is not unheard of.

It will send to someone at Berkeley. The Human Resources/Payroll Department has completed the final paystub blah, blah, blah. ADP PORTAL claiming to come from ADP PORTAL, which is used in the U.S. by many companies for payment purposes for staff and pension-related matters and all of that stuff.

So, this was easy to catch because it is evidently coming from a different domain. This is just a compromised server but it's much harder to detect when it comes from a very good look-alike domain. So, it does exist and it does happen.

And as I was saying, the response to the gentleman's question a little before, there are some techniques that can be used – technological solutions better that can be used by people and companies to mitigate certain types of threats. Our team is always on the road. We're always out in many countries training law enforcement, speaking with the government officials. And in general, people within their public safety community.

This is something that I used to highlight because it's important at least the way I see it, say the financial industry or the government entities within a country.

They can implement these measures that I implement, that I mentioned, which are too technical to discuss right now but at least I'll just mention them and just say their names so that you know them or at least can tell that you've heard of them. SPF, DKIM and DMARC – in a very, very 30-second way to put them, they allow for domains not to be spoofed by criminals.

So if BankofAmerica.com implements these three measures, the result will be that no e-mail coming from BankofAmerica.com but having been sent by the criminals, meaning criminals spoofing BankofAmerica.com, it won't go through. And the real guys at Bank of America will get a report and will know that someone is trying to use their domain for malicious purposes.

It is the DNS resources that's why I talk about it. It has to do with including certain information in the zone file for the domain names that are being for which this is being implemented. And that information is spread throughout the entire DNS, and mail services, mail service providers and the mail clients, all use that information, the end result is to protect users basically.

So, it's just good Internet citizenship I guess. It won't prevent the companies from being attacked, it will prevent the population. That's the whole thing.

So, learning more about abuse, you don't need to write these URLs down. There's cool information here from MIT, Malware Domains, Ransomeware tracker, SANS Storm, Internet Storm Center – this is always a good source of information.

There are many, many, many, many sources out there. If you look up DNS abuse, you're going to find so much information that you're not going to know where to start, so just start somewhere. Pick something, pick malware, pick botnets, pick command and control or Domain Generation Algorithm and just get started. It's very interesting.

If you have any question about this, please let me know. You can reach me at Carlos.Alvarez@ICANN.org. You can just write to me. That's why we're here for. We're here to assist you. We're here to help you go through these topics and to help you understand them. If you are a law enforcement or in the security community or come across a malicious domain and want to know what you can do with it or what someone can do with it, we can help you understand what can be done or what could not be done.

So, thank you so much. I think we're almost in time.

STEVE CONTE: Yeah. We do have some space for questions if anyone has any questions, go and raise your hand. No.

CARLOS ALVAREZ: Please, yeah. That gentleman.

STEVE CONTE: Please.

UNIDENTIFIED MALE: Thank you [inaudible] for detailed presentation on DNS abuse. My question is that you mentioned that the point of contact for the domain DNS abuse is registrar. Sometimes the registrar is not working efficiently and sometimes the registrar is basically the registry operators. So, at that time, what is the role of ICANN? How the [affectee] will report to the ICANN?

CARLOS ALVAREZ: In general, and out of what I had seen out of being in this environment for sometime, registrars will act when [inaudible] reports of abuse that have to do with phishing, botnet, command and control, and malware distribution, and child abuse material as well.

Ideally, you would provide them with as much information as you can up front. If you're talking about a malware, a domain that's used for distributing malware, not that you necessarily have to but it would be good if you provided a copy of the malware sample and some trace, crowd information, so just some basic information that can allow the registrar to take action.

As I said, now that you have to but it's better because you're going to help the registrar to take action more efficiently. If you feel that the registrar doesn't take action promptly or if you feel that the registrar didn't take reasonable action, you can go to the Compliance Team. The Compliance Team – you can look it up. What they do is I don't go directly to ICANN's homepage. I just look up ICANN Compliance complaints. And among the first links, you'll be able to go to the complaint forms, look under registrar abuse and you'll find a form that's called registrar standards or abuse something, and there is where we can submit the complaints for Compliance to follow-up on these types of matters, that's basically the way to go about these issues basically.

Does anyone else have any other question?

STEVE CONTE: All right, thank you Carlos for this.

CARLOS ALVAREZ: Okay.

STEVE CONTE: I really appreciate your time here.

CARLOS ALVAREZ: Thank you.

STEVE CONTE: We're going to take about a 15-minute break. And, coming up at 3:15 in this room, we'll have some root operators from the Root Server System Advisory Committee. They're talking about the Root Server System –

[END OF TRANSCRIPTION]