KOBE – How It Works: Understanding DNS Abuse
Sunday, March 10, 2019 – 17:00 to 18:30 JST
ICANN64 | Kobe, Japan

CATHY PETERSEN: Good afternoon, everyone. Welcome to our How It Works Tutorial on Understanding DNS Abuse. Today we will have a new presenter. This is Samaneh on my right. She is from ICANN Office of the CTO. Thank you.

SAMANEH TAJALIZADEHKOOB: Hi, everybody. Today we are going to discuss some general concepts about DNS abuse. Feel free to ask questions whenever you feel like you don't understand anything, etc. So, this is the outline of the presentation. First we will discuss what is DNS abuse in general and how is it different from DNS misuse. Then we give some examples of DNS abuse, then we look at the threats landscape and then we finish with what is ICANN doing in terms of abuse.

So, the first and most important point about DNS abuse is that there is no very specific definition or a definition that everybody agrees on. Some people define it as cybercrime. Other communities call it hacking, kind of malicious conduct. But what more or less everybody agrees on is that DNS abuse involves data corruption, denial of service, or violation of privacy.

So, how is DNS abuse different from misuse? Let's say DNS abuse refers to when attacker abuses the DNS infrastructure, whereas DNS misuse is exploiting the protocols that are used in DNS system for malicious purposes.

Why is DNS interesting for attackers and why is it targeted of many attacks? Well, the basis that – this is the base of the Internet. If you want to browse to websites, you need DNS to resolve domain names to Internet protocol addresses, the IP addresses. So, attacking the system means disrupting websites, businesses, e-learning platforms, social platforms and many other businesses that is using Internet as base.

Possible or currently used [vectors] for exploiting DNS involves registering domain names or compromising servers to use domain names for attacks, hijacking name resolution, or registration services, and in one way or another corrupting DNS zone data to perform attacks.

Before going into more details of the type of attacks, let's take a quick look at how DNS works. So, to put simply, a client – let's say, you – want to browse to Google.com. Your machine – [inaudible] machine – asks resolver or basically queries the resolver about … Yeah, queries the resolver for the domain name and the resolver goes to authoritative servers to basically ask, "Where can I find this domain?" So, the client or the stub resolver is one part of the DNS resolution that is in between and the authoritative server.

In this process, if the client or the stub resolver caches the answer, that would be for the time to leave. So, the TTL of the cache. That would already give the answer to the client or it would go to the recursive resolver. And if not, at the end, it will go to the authoritative name server which holds the zone for the TLD of the queried domain.

If we look at what elements of the DNS infrastructure is targeted, we see that there are different, both software and hardware that is targeted. Some of them require bandwidth, some of them has OS, has cache, and of course, if it involves browser, it will have application software, etc.

Then we move to the examples of different attacks that are performed. This is just a list. This is not all of the possible attack vectors. We will look into detail of these examples one by one.

I'm sure most of you have heard about distributed denial-of-service attack, or DDoS. This is a kind of attack in which the attacker overloads the IP address of a target without the target machine asked for the traffic. So what happens in this attack is that attacker uses a spoofed IP. Let's say in this example, a target IP 10.0.0.1 and uses an open resolver that is a badly maintained, or badly configured to send a large amount of traffic to the target IP.

There are several open resolvers available on Internet, operated by Google for instance, or a Cloudfare. Those are the well-maintained ones, but there are also open resolvers that are not so well maintained. And those are the ones that attackers could use to perform such

attacks, sending UDP packages. In this case of reflection and amplification DDoS. In this kind of attack, the receiver of the attack or the target can receive ten or a 1,000 megs of traffic on its machine, which can completely put a server out of service for a while.

The resource depletion DOS attack is also kind of a similar attack, with a difference that instead of sending UDP packages, the attacker in this case use TCP packages, until the resources of the receiver are completely exhausted, because the name server of course has to allocate resources for the TCP traffic that is coming towards the server.

What are the possible solutions or defenses for such attacks? First of all, providers, or most of the providers, currently have solutions, like mix of provisioning or filtering techniques that they use. It's recommended to use software that is up to date, and commonly used. And basically most of the protection on this kind of service should be done on the upstream, as in the Internet service provider, or the cloud provider. So it's hard to be done by the client itself.

Then we move on to cache poisoning. In cache poisoning, basically attacker launches campaign to direct [inaudible] … Basically for you to browse to a attacker controlled website. In this case, loseweightfastnow.com. And as soon as your browser tries to query the example website, the attacker – the [inaudible] controlled by the attacker, answers with a IP address that belongs to the attacker, basically a malicious IP address, that points to the server that attacker

controls. So what happens is that the resolver caches the answer. In this case, the resolver wrongly points to ww.ebay.com, which is not currently controlled by the attacker.

So what happens is that the browser caches this response, and for the time to leave the TTL. And whenever later, depending on time to leave, client want to browse through this website again, it will redirect to the same malicious server, which then can be a source for different other, further types of attack, like downloading malware, harvesting credentials, etc.

Possible solutions for cache poisoning involve, first of all, keeping resolvers private to yourself. So, not let the resolvers be exposed to public. Use patched and up-to-date DNS software. There are a lot of studies that show a big majority of DNS servers online are surprisingly not using up-to-date software and are not fully patched.

There are some software that use port randomization techniques, square IDs, etc. for avoiding cache poisoning. So that is also something that providers would like to look into. And DNSSEC would be also a technique to validate queries and responses, signing and knowing that this is a valid response you are receiving, and not from the attacker.

Another type of DNS attack is poisoning [inaudible]. So, I already kind of explained this. That when, for instance, your machine is hacked, or it can be that your - yeah, you have a malware on your website that you got through different means, through clicking on a spam link, etc. We

don't go there. When you have a DNS configuration malware installed on your machine, it's also called a DNS changer. What the malware does, is that it alters the DNS configuration of your machine.

Every machine has a DNS, basically, server that is already configured in the machine. What this malware does is changes this server to a server that is controlled by the attacker. So, what happens is that every time the client queries, performs NS queries, it will go to basically a destination of attacker’s choice, which would then be a starting point for other types of attack.

Again, this is more of a network hygiene problem. Meaning that in the first place, the client machine is compromised. So, it's good to look for not clicking on spam emails, or basically not getting hacked or downloading malware. Keeping software patched, and most importantly, monitoring DNS traffic for anomalies or unusual behavior.

What is also important for network hygiene kind of problems is user awareness or if it is a company staff training about what kind of emails look like phishing. Once a while, try your staff to see how good they are in detection of phishing emails, not clicking on links that download malware, things like that.

There are a bunch of attacks that are more on DNS protocol and registration system. This is a list and these are, again, just examples of possible attacks.

In general, attackers … So, as I explained in the previous slides, domain names are used to carry out all kinds of DNS-related attacks and they are used in mainly two types of categories by attackers. Either attackers register domain names to use them from the start for distributing malware or other types of things, botnet command-and-control servers, or they compromise machines to take over already existing domains.

For instance, in the case of shared hosting, which is a situation where several domain names are hosted in one server or one machine, if an attacker manages to get hold of vulnerability in the server, which is quite common, then it's really easy. Once one domain is compromised, then it's easy to take over a group of domains, which is way more cost efficient than registering domain names from scratch. But it's also very common to register domain names in bulk for malicious purposes.

Such purposes involve phishing, pay a ransom payment webpages, counterfeit goods, illegal or pharmaceutical websites, piracy sites, and all kinds of child abuse material. It can also be used for name servers, and like I explained, command-and-control administration of botnets.

Another type of attack is hijacking or theft type of attack. In this case, the attacker takes control of a domain registrar or registry, and uses – or customer account, basically harvest credentials somehow, or first compromises the customer account and then get access to the credentials. And [users] start through social engineering, phishing attacks, data breaches, etc. And then, once the attacker has control

over the account, then it's really easy to change the NS server of the domain and points it to the server that the attacker is controlling.

And then, let's that in the zone file, and let the zone file be published. So be announced to the – basically the whole DNS ecosystem, which then means that if somebody queries that domain later, it will point to … The response will involve the NS server that is controlled by the attacker. There is also a lot of attacks being carried on – compromising, exploit registrar emails, correspondence. So, basically, again, taking over accounts, credentials, etc. In the past years we have seen several examples of that. The Yahoo or Equifax breaches and a lot of other examples.

Possible solutions for not getting victim of such attacks involve, most importantly, two-factor authentication. So even if the machine or the device of the customer that is used to log into the registry or registrar platform is compromised, a second mean of authentication could be a backup option as a defense mechanism.

Name server locking and registrar and registry locking are also very important mechanisms against such attack and a DNSSEC, as I already mentioned, could also be helpful in this case.

It's good to have support, dedicated software staff, escalation path, and clear anti-abuse policies in case abuse happens in such and such scenarios. Then we see … Before I move forward, is there any questions?

CATHY PETERSEN: Yes, do you want to – sorry. Do you mind using the microphone over here? Thank you. And kindly provide your name and affiliation, please.

NARELLE CLARK: Narelle Clark, I'm with PIR's board, and I'm also a member of auDA and involved with auDA's .au policy panels and so forth. Question I just had was, quickly, do you want to give us a run down a bit on the recent DNS hijacking if you could, and if you have time?

SAMANEH TAJALIZADEHKOOB: It's not part of this presentation, but I could talk about it after this session with you if you are interested, and if time allows after this presentation in the public session. Any other question? Okay.

So, I move forward to basically fast flux – attacks that use fast flux technique to carry on DNS-related attacks. What attacker does, is that they associate IP address to a proxy server or a name server for a short time to leave. And this proxy, or this IP address, keeps changing for the proxy or the name server. And this is kind of a technique to basically evade detection, or if one server goes down, there will be back-up option mostly for attacker infrastructure, like command-and-control servers, that are critical to be up all the time. There's also double fast flux that applies this technique with more … To both proxy servers and name servers. DNS can also be … yeah, question?

RON ANDRUFF: The double fast flux and the fast flux have been around for about ten years now roughly. Have you seen that that has become the standard in terms of how these are being used or has it increased or decreased in any sense? Because it's such a powerful thing in terms of how they've been using it, and that's how they stay one step ahead. I guess, are we able to catch them at all? Is there any white hat stuff going on in that black hat scenario?

SAMANEH TAJALIZADEHKOOB: Thank you for the question. There is certainly detection going on, but because of the technique, it … I mean, the technique itself also makes it hard to be detected. Short answer is that I'm not sure if it is decreased or increased, but I know that it still exists pretty much. So, the current defense mechanism are not able to fully detect it.

RON ANDROFF: Yeah, that was what I was asking. Thank you very much.

SAMANEH TAJALIZADEHKOOB: You're welcome. Any other questions? Okay. So, DNS is also used as an exfiltration channel. This means that DNS messages can also be manipulated to forward sensitive data from the victim's machine and bypassing my firewall to attacker – a machine that is

controlled by the attacker. For instance, it's very common in command-and-control servers. The attackers use port 53, which is an open outbound port for the firewall to communicate. Basically, to transfer sensitive data from the machine or communicate with a bot to the botmaster or the command-and-control server.

It's also used as a malware channel, meaning that the command-and-control machine uses inbound – basically port 53, that is open inbound also. And when the attacker asks for text record to the botnet CNC, the CNC sends instructions via the test – attached to the text records to the infected machine. And you see this example in feederbot or morto. You could look up the details from the source in this slide also.

How is the DNS threat landscape evolving? Well, what we see is that more and more DDoS … Performing the DDoS attack is becoming easier. DDoS is offered as a service and so is spam is offered via cloud. It is offered via messaging, social media, etc. Examples of such are Avalanche and Avalanche and malware and DNS infrastructure, hosting infrastructure. And then, there is Internet of Things that took botnet recruitment basically to the next level. We are connecting more and more devices to the Internet that are not very well thought through in terms of security. And we are offering them to users and motivate them for buying insecure devices through cheap prices, etc. So, basically, there is nothing that backs up such devices.

Going more specifically to the case of DDoS, currently DDoS is offered as a service. You can hire boot, the so-called booters or actors that perform DDoS for you as a service – booters or stressors. There is a large body of academic work on this topic, and you see that such service can even be offered as cheap as $1 per DDoS attack, depending on the amount of traffic one wants to send. So, it's something that is really easy to access, for people to use it.

There is DNS … Basically, malware on DNS hosting infrastructure. I'm not sure if you have heard about Avalanche malware. Basically Avalanche evolved from being a botnet to malware delivery service. It used bulletproof hosting to host names and double fast flux technique, and it's mainly used for financial fraud attacks. Basically, this was the first malware that offered cloud customer experience, everything on cloud. Criminal domain registration, access to CNC servers. There were 20 malware families, so the person used the service, had choice of different malware families to use.

This is how the malware evolved. Basically, it started from a certain botnet and evolved to Rockfish botnet and evolved to Avalanche, evolved to being a malware delivery service, and then criminal investigation started to pick up on it, in 2012 ad in 2016, they took down the botnet, and had arrests and other [inaudible] investigation going on.

The attackers, as always in this field, were very proactive while the law enforcement was catching up with them. It took a while before the law enforcement was successful, but it was an impressive outcome.

There were five arrests in four countries. Several [searchers] in seven countries, 64 TLDs involved. And they were also ever … As a result of the campaign, there were awareness raising about the malware, and the possible prevention techniques.

But did it end there? No. Like a lot of other malware campaigns, this was the start for a new type of family. The source code was used to develop Andromeda, which was a direct follow up of Avalanche.

Like I said, this is something that is repeated and repeated in the history that attackers use – new and custom OSs, software apps to develop new malware based on the current existing code, new actors. We also keep falling for the same previous mistakes that we already experienced.

Mirai is another example of a basically botnet malware that used IOT devices as a botnet or bots. What happens is that the malware scanned for using default credentials of the IOT devices which are a lot of times public. So, if you already have an IOT device that is using the default username or password, be aware that it can also be public. So, the malware used this public default password to compromise or take over the IOT devices and use these devices as bots performing other types of attack. And because these vulnerable IOT devices are typically listening

on port TCP23 or 2323, it was really easy to compromise these devices through a basic text-based communication channel.

This was a really interesting point in history that we basically added a lot of more vulnerable devices to already existing infrastructure of Internet, made it easier for the attackers to create botnets. And machines that can basically be used for all kinds of attacks within a home base of users.

What this means is that this is a base for … Basically this took DDoS as a service to a new level because we are using machines of customer themselves to perform DDoS.

There was also WannaCry or WannaCrypt ransomware that was based on the vulnerability in Microsoft OS. The vulnerability was exploited three months after announcement, which was not patched. It came in several waves, and at the end, it was taken out by researchers [sinkholing] an unregistered [CNC] domain.

Having all of these examples in mind about DNS attack, I want to put now what is explained in the context of ICANN's role about DNS. ICANN often touches on the concept of DNS and takes the role of basically helping registries and registrars to clean their networks of abuse. Current topics of highest priority in ICANN 64, and that is related to DNS abuse, involves WHOIS accuracy, GDPR that is related to WHOIS, public safety and abuse reporting.

The current project within the OCTO, our department, that currently I'm working on, and some of you may have heard of it, is called DAAR, the Domain Abuse Activity Reporting. And that involves collecting several abuse feeds, and reporting these abuse concentrations for TLD registries and later for registrars, and if possible ccTLDs. Later this week, there will be a session on this I think on Wednesday morning. For those who are interested, can follow up the work there.

The discussion of abuse basically also went to the GAC, the Government Advisory Committee on DNS abuse. It is discussed that what kind of abuse are important from the point of view off GAC for domain names, and basically what is important to be focused on when looking at abusing your network as a TLD provider or registrant.

As a result of this discussion in GAC, there is a Public Safety Working Group formed. This working group reports and advises GAC on matters related to abuse and public safety or public interest policy, works with law enforcement and cybersecurity SMEs and issues the PSWG about GDPR, WHOIS accuracy, carrier grade network, address translation, fast flux, DNS abuse, among other topics.

As a result of this, there was also a consideration of DNS abuse in contractual agreements with registries and registrars. Basically, there were new specifications added related to the fact that providers should be aware of abusing their networks. So, monitor their abuse, and

basically have methods of dealing with it, mitigating it, and reducing it actively.

That's it for today. If you have questions that you want to discuss through email, this is our email address. John Crain, the Chief Security Manager, and I the Security Specialist in the group. And I'm happy to answer questions, if there are any now.

CATHY PETERSEN: Any questions or comments? No? All right, I think we're going to end this session early. Give everybody some time. So, thank you again for attending our session. The slides have already uploaded in the public schedule and the link to this recording will also be posted on the public schedule within a week. Thank you again, everyone.

[END OF TRANSCRIPTION]