BARCELONA – Tech Day (1 of 4)
Monday, October 22, 2018 – 09:30 to 10:15 CEST
ICANN63 | Barcelona, Spain

EBERHARD LISSE: Okay, good morning, everybody. Please find your seats. As you know or not, my name is Eberhard Lisse. I'm the ccTLD manager of dot-NA. I usually make a joke that in my day job I'm a gynecologist, and in my night job, I'm an obstetrician but I'm not doing the night job anymore because my practice insurance.

So, this is the 38th iteration of Tech Day, so we have been doing this for a number of times. This time, we also have a good number of presentations. We ran into the problem of having to remove the coffee breaks because we had enough presentations and we even had one in abeyance that we only could put on the day before the meeting before somebody – well, somebody canceled on short notice. I personally think this is a good problem to have. What we usually do …

Some housekeeping announcements beforehand. We have got a boxed lunch. Only people that we recognize having been here the whole morning will get a ticket at the door. Without a ticket, you won't get a box. The box is sponsored by SerNet. Johannes Loxen from SerNet will have the presentation before lunch, as I think is fitting. If you want to speak to him about anything, then the lunchtime would be an opportune time to do this and join him at his table. Good.

What I usually do is I quickly go through this, if it was working. There you go. Are you having me … Or is it a manual clicker that I push the button? Is it working? Okay. Alright.

First, will be a not directly DNS-related topic about bug bounties, but I think we found that this is something that might be interesting for anybody just to get some concepts on how to deal with this.

Then, we have the presentation from dotCAT. dotCAT is not a host, but I found it appropriate since we are guests here to invite them, too.

Then, we will have a presentation about the [inaudible] of the GDPR on zone poisoning. The exact title of the presentations may differ from what we see here on the agenda because it has to fit on one line and I even had to reduce the font size anyway.

Then, something about Anycast, peering, and sinkholes. Then we will have a demonstration, a recorded demonstration, by Jacques Latour about the secure home gateway where they use DNS to secure the Internet of Things in their houses. Then, Johannes Loxen I understand is speaking about Next-Gen Firewalls and some issues that they noticed when they deployed them.

Then, we have I think the most interesting presentation. The people from dot-AU are going to tell us how this works. This is the biggest TLD transferred and largest number of domain names transferred and I would like to know how this happened. What were the pitfalls? What were the issues that they encountered? Why not, being interested in too much of the details, it's surely going to be interesting.

Then, dot-AS is going to make the host presentation. Like dotCAT, the host presentation we usually offer the host a presentation of their choice and also would like to hear a little bit about how that is set up, so that we have an idea of what's going on, but also that they can present any topic of their choice.

We have had RDAP updates from Francisco Arias before, so he asked us to give us another one. Then, Brett Carr from Nominet is going to talk about Anycast in the Cloud.

Then, the topic that I find personally very interesting is open source HSM. I find that DNSSEC is very easy to do, but very expensive to do in an auditable way and very difficult to do in a secure way. If there wasn't open source HSM that was auditable, that will be [very cool].

Then, SSAC is going to talk about the newly noticed IDN homograph attack. I was explained yesterday what that means but I don't have to explain it to you.

Then, Geoff Huston is going to give us measurements about the KSK roll. I personally didn't notice anything other than that for a few days I had intermittent speed issues with my connectivity overseas. Maybe it has something to do. Maybe it was just something incidental, but even in Africa we didn't notice much, probably because I use the Google name servers and they are not affected by this.

Finally, [inaudible] is going to talk about the WHOIS use at dot-NO post GDPR, the changes that is required. I wanted to put it together with the other GDPR presentation but due to a late schedule change, we had to

move this apart. It's a bit of a nuisance, I think, that we have to split topics because people who want to, for example, are interested in the technical part of GDPR could come to this block of things, but it couldn't be done.

Then, we will find somebody, and my eye is falling on [Jao Mitolio] from dot-CZ who will give us the summary. We usually do that one, other than the chair, is going to give his or her own view of what he heard and what she thinks about what he heard. If he's available, I will ask him. If he's not, I will find somebody to do it.

That said, Gavin Brown has the floor.

GAVIN BROWN: Thank you, Eberhard. Can I have the presenter? Thank you. My presentation today is on running a bug bounty program for a domain registry or registrar. Just to give you an introduction, I represent CentralNic. If you haven't heard of us before, we're a group company that contains four different registries and five registrars. So, the registries include CentralNic, the registry that runs dot-XYZ and other top-level domains. OpenRegsitry, KSRegistry who came to us when we merged with Key Systems early this year, and SK-NIC the ccTLD operator for dot-SK which is Slovakia.

We have a number of registrars, TLD registrar solutions Internet-bs, Instra Corporation, and Key Systems and obviously Key Systems includes registries like RRPproxy and Moniker.

We have ISO 27001 certification for several parts of our business and a general strong focus on security, especially in our registries. Obviously, running critical infrastructure means that we have to take that very seriously.

As operators, we have bug bounty problems for the three of the companies within our group. CentralNic, Instra, and Internet-bs which was essentially the company before we completed the acquisition of SK-NIC and Key Systems.

So, what is a bug bounty program? To explain it in one sentence, a bug bounty program is a continuous crowdsourced black box penetration test. You may be familiar with traditional penetration testing where you go to a security firm or consultancy firm, give an enormous amount of money. They go away and come back a few weeks later with a big, thick report listing all of the vulnerabilities they found in your software when looking at it from the outside.

This is a different approach to the same problem. Instead of engaging a single company, you basically invite the whole world to come and hack you. Independent security researchers will test your systems. Sometimes, they'll use automated tools. Sometimes, they'll use their own ingenuity. Find problems with it, submit reports, in the expectation of receiving money for the vulnerabilities that they find.

A number of large companies use this approach for doing vulnerability detection. Some of them are in this room, companies like Google, Facebook, and the list is there you can see. Often, these programs are managed for you by third parties. There are a few good ones.

HackerOne is the company that CentralNic uses. Bugcrowd is another. There are others as well.

A little bit of information about CentralNic's bug bounty program. We've been running it since 2015. We use the HackerOne platform, as I mentioned before. At the time that this slide written late last week, we had 451 hackers participating in our security program. The vast majority of them don't really do much. Most of our reports come from a much smaller number, 36 in total. We have received 345 reports over the course of those four years – three years, sorry – of which 206 were legitimate. So, reasonable fraction. I think if you've ever had a traditional penetration test, you'll be aware of the false positive rate, especially if automated scanning tools are used, where the vast majority of the data that is produced from that report is bogus and needs to be manually confirmed and verified. One of the nice things about a bug bounty program is that someone else does all that work for you and you don't have to pay them for it.

A reasonable number are also out of scope or simply had no security impact, reporting something that was simply not a security issue at all.

A smaller number were what the HackerOne system calls informative. That was to say it told us a little bit that we didn't necessarily know about the way our system looks to the outside world. It didn't have a security impact that required us to fix, but it was useful when thinking about security considerations in the future.

We've had a few duplicates where two people were looking at the same thing at the same time and one of them submits a report shortly after

another one. We have to say, "Sorry, this has already been reported and you'll have to just accept the fact that someone got there first."

One of the nice things about a managed bug bounty program is that we've had zero spam. We have never had someone just sending junk through the system because these managed systems, they filter that out for you and they're also heavily reputation based. Our bug bounty program is private, so you can't see it if you're not a member of HackerOne. But the security researcher's profiles are public. Their signal score, their reputation, is visible to the outside world. I'll talk a little bit later about why that is and how it generates favorable conditions for doing security management.

So, a few more stats about our bug bounty program. We've paid out over $20,000 in bounties. One thing to note is that wasn't all paid out in one go, in one lump sum. That was paid out a few hundred dollars at a time over the course of several years. So, from the point of view of the cost of this versus traditional penetration testing, I think it looks very good because $20,000 is easily the kind of money you'd spend on a traditional one-off penetration test by a security company. But, of course, as soon as that document is written, it's no longer valid because if you're doing continuous deployment and continuous improvement, then the next time you push an update to your production environment, then that report becomes invalid.

We pay around about $100. Now, that's very small compared to some of the largest bug bounty programs who pay thousands of dollars even for very minor vulnerabilities, but then they're protecting hundreds of

millions, if not billions, of users. It's proportional to the impact of those reports.

As you can see, 206 reports resolved and 36-37 hackers involved who have been publicly thanked by us for finding issues.

The third-party service that we use, HackerOne, implements SLAs on the companies participating, the companies running the programs, as a way of giving the hackers the visibility on who are the kind of people you might want to work with because they're the ones who are most responsive. So, we try very hard to make sure that our number look good. As you can see, if you submit a report to us, we usually get back to you within 24 hours. We usually triage it within the same sort of timeframe and we can usually resolve an issue and pay out a bounty within around a week, and obviously we're doing quite well in terms of meeting our SLAs.

So, what sort of reports do we get? Now, I think this is probably quite applicable to most companies. Quite often, these are sort of things that are oversights rather than serious flaws. It's things like TLS configuration. So, if you're standing up a web server or an SMTP server or a JAVA server where you're running a TLS port, if you're just using an out-of-the box configuration, that's not [inaudible] and the algorithms that you're using, the likelihood is that a weak cycle will get in there somewhere. And obviously this is a moving target because the [cyphers] and the algorithms that are known to be safe changes over time as crypto-analysis comes along and finds problems with existing [cyphers].

Session management is a very common one. Things like lack of rate limiting on log-in forms to defeat brut force attacks. Things like not logging other sessions out when you change your password. So, if an account has been compromised and a user goes into the system and changes their password, that should invalidate all of the other log-in sessions associated with that user account, because obviously if you're trying to prevent someone from getting access to your system and they've already got access, you want to knock them out.

Those kinds of things which are enhancements over and above the simple. Is the password correct? You need to think about. And the bug bounty programs are very good at encouraging best practice because each time they find one of these best practices you're not following, you'll get a report and then you can fix it.

Things like hardening headers. The content security policy and the XXS protection. There's a website called SecurityHeaders.com which lets you test your web server to see if you've deployed all these headers. Obviously, they're very easy to automate in terms of testing. So, someone can just take your /24 and check every server to make sure you've got the correct headers in place, and if you haven't, then they'll submit a report.

So, this is, again, a good way of improving the use of best practice across your organization, across your operations, because these are the sort of things that will be picked up by security researchers. It's cheap for them to find. It's cheap for you to fix. And then you're paying out a

relatively small bounty, but you're still improving the security of your system.

There are a few things that I find quite annoying but we still want to keep in scope. Things like information disclosure. So, that's what version of a patch are you running? A question of what PHP are you running? The sort of things that leak naturally out of a lot of open source software, which you may not necessarily want people to know about because I know you're running an insecure version of a patch or an insecure version of PHP. That makes it easier for me to get into your system. So, you may want to start removing those version numbers from your error pages and that kind of thing.

The final one is text injection. This is, again, a bit of an annoyance of mine. It can only really be exploited in conjunction with something like a social engineering or a phishing attack where someone finds a URL that takes a text field. It's being properly escaped, so you're not vulnerable to across site scripting, but it's saying, "Please go to BadSite.com instead of this page," and now your user might be convinced to follow that link, essentially, and type that site into their browser and then be taken off somewhere bad.

So, what are the benefits and drawbacks of running a bug bounty program? As I mentioned before, the main reason we did it was because it was better on the bank balance over the long run. Over the course of three years, we may have spent the same amount as a penetration test, a traditional penetration test, on paying out bounties, but we only paid it in small amounts over a long period of time. And if you are a small

country code TLD or a small registrar who doesn't necessarily have a security budget big enough to accommodate regular penetration tests, then this can be a good solution, a good alternative solution, to a traditional penetration test.

The other thing is that it's continuous. As I said before, if you pay for a traditional pen test and you get a big 500-page document telling you all about all the vulnerabilities in your system, as soon as you push an update to your system, then that document becomes less relevant and less accurate, and as time goes on, it becomes less and less accurate and then you get into a situation where you have to repeatedly test your system at great expense in order to pick up any new issues that may have been introduced during the time since the report was published.

A bug bounty program resolves that issue for you because people, they will be continuously testing your system as they enter and leave the program, as they take their interest in your program's waxes and wanes. They will be testing stuff that previous people have tested before.

Another thing. The last two points on this list can be considered together. It generates good will. There is a security community out there, a hacker community out there. They are looking at your systems anyway, even if they are not telling you they are. And this is a nudge. It says to them … You have to accept the fact that they are looking at your systems anyway. Here is a little nudge to go from being black hat to white. It keeps the honest people honest and maybe the less honest

people are inclined to become more honest because they get a benefit in a monetary value and a reputational value.

So, what are the disadvantages? Well, the first one is the obvious one, and in this reasoning, which you're airing your dirty laundry in public. You are saying, "Here is my stuff. Try and steal it," which scares a lot of people, especially those who don't necessarily have the right mindset to think about the fact that this stuff is all already visible anyway, especially if you're running infrastructure. If you're running a registry or a registrar, it's all out there anyway.

It paints a target on your business because not only are you saying, not only is your infrastructure out there, you're actually inviting people, and that again is scary.

One real concrete problem that we've come across is that when it comes to third-party potential vendors – or sorry, potential clients – looking at your business from their own security point of view, not having one of those big 500-page reports that you can slam down on the desk, it does change how they see your business, because quite often, they'll come to you and say, "We want to see a security report. What was the result of your last pen test?" And we can't say to them, "Here is the result of our last pen test." Slam. What we can say is that we have a security research program which is a continuous process and here are some stats about the rate at which we identify reports, the rate at which they're resolved. But, it doesn't quite have the same impact as a great big document with a well-known security company's logo on the front.

The final thing is that it does require quite careful management in order to be successful. You have to have commitment within your organization and you have to – and the subsequent slide will show you do need good internal procedures in order to handle the reports as they come through.

So, let's say you want to start your own program. What do you do? How do you go about it?

So, the first thing to do is find a provider that works for you. Start with a small scope. Look at other security programs. Look at the exclusions that they've come up with. So, don't look at this particular type of vulnerability, don't look at this particular system, and think about how you could adapt that to keep your scope small and then aim to gradually broaden the scope as you get more and more confident about the security capabilities of your system.

Make your bounty calculation rules consistent and transparent and fair, because if it seems like you're being arbitrary with the way you pay out bounties, people won't participate. Use incentives to direct hackers to particular areas. It's very easy to do web application and vulnerability scanning because there are thousands of tools, hundreds of tools, out there that can do it. If you want them to target WHOIS or you want to target EPP, then you maybe need to use a monetary incentive – a bonus. If you find a vulnerability in my EPP server, you get twice as much money as you would if you found a similar vulnerability in a web application.

Make sure that you have an ultimate aim to go public where anyone can submit reports, not just invited hackers. And make sure you can provide test accounts. When I say test accounts, that may mean on [OT&E] but only if your [OT&E] environment is exactly 100% reproductive of production. If not, then on production.

How to deal with hackers. These people are self-employed security professionals. They might use the word hacker, but they are self-employed security professionals. They're often quite new in their careers. They will try and use their experience through the bug bounty programs to leverage it into jobs in security world. We've had a situation where people participating in our bug bounty program have actually gone on to be employed by HackerOne themselves. So, that's a real thing that happens. Most of them don't speak English. They are not always elite hackers. They might just be starting out. So, you be patient and you be polite.

The false positive rate will always be lower than automated pen testing, but you will have to deal with noise. So, ask for evidence. Ask for proof of concept. Set yourself SLAs to provide quick responses. And as I say, you must back your program onto robust change management and Q&A processes. If you can't get fixes into production quickly, then there's no point in running a bug bounty program because you'll just keep on having to close duplicate reports because you can't get the changes out there quickly.

Also, make sure that you've got good operational practice in terms of configuration management, so you can fix once and deploy everywhere.

Calculating bounties. We have a simple system. It's just a product of two numbers and then a sliding scale. Again, it's about being able to justify your calculation, having an objective number. Even if that formula ends up paying below the market rate for that sort of vulnerability, as long as you're consistent and as long as you're transparent, that will be fine.

What to expect if you run your own program. Some spam, some false positives, but still lower. Okay, we're done. That was the last slide anyway.

Don't forget you can tune your program to improve the [inaudible] ratio and make sure that you've got good change management internally.

So, that's the end of the presentation. If you have any more questions, we can be reached at these links. Any questions?

EBERHARD LISSE: Thank you very much. One thing I forgot to say, all the presentations plus the agenda will be posted. The agenda has an e-mail link, so if you want to contact the presenter or one of the presenters, just click on the link and your mail software should pop up and you can then ask questions. I can allow one question. I see in the back we have got standing room only which is another nice problem to have, but it's a few seats spread around that you can still sit down. I can allow one question

because we are running a little bit behind time. And when you are asking, please identify yourself because we have a remote audience.

UNIDENTIFIED MALE: [inaudible]. How did you protect your customers from these hackers? What if the hacker disrupt your service or register maybe a domain or took a domain from [inaudible]? What you will do in that case?

GAVIN BROWN: So, there's no technical way of doing that. You have to rely on the program policy and ask the hackers to obey that policy. Obviously, if they don't obey that policy, then you are under no obligation to pay them a bounty, and [inaudible], you get a free vulnerability report. If you're concerned about that kind of thing happening in reality, then the solution is [OT&E]. Use your test environment. Provide a test account. We are lucky. CentralNic is lucky that obviously we run many, many TLDs on the single platform, so we created a test TLD in our system and gave the registrar access to that TLD only. Obviously, if they found a particularly bad issue, a bad vulnerability, then they can bypass that access control. But that's another way of solving that problem. That's my answer.

EBERHARD LISSE: That was the answer to a question that I was having in my mind. It's [OT&E] on a separate system always runs a chance of not having exactly the same versions, so putting a separate test TLD into the production

system is a good idea. Okay. Thank you very much. I like this presentation quite a bit. Give him a good hand. [applause]

Next we have Pep Masoliver and Nacho Amadoz from dotCAT. I'm very happy that he arrived today because he has a fascinating hobby. He is a part of a team that builds these human castles that are of cultural importance. He invited me yesterday but it was shortly before I had an appointment which I couldn't break, so I'm really upset about it. But we'll see something whether there is something next Sunday because this is something that I'm going to see for sure.

PEP MASOLIVER: [inaudible]. First, thank you for the information. Human towers are simply a Catalonian cultural characteristic. We will try to present the domain and one of our changes that is a privacy by default since the very beginning.

As you probably know, the dotCAT represents the Catalan-speaking community which is about 13 million inhabitants and 10 million of them are habitual or potential Catalan speakers.

The dotCAT domain is managed by the Fundacio puntCAT which is a non-profit NGO.

The dotCAT domain was approved on 2005 and we launched it on 2006. Now we have around 100,000 domains. And we are only 12 people, 8 full-time and another 4 part-time.

The technical team is focused on innovation and support. It is not … Because we have an external provider for our background, which is CORE. You may know them. We are using their TANGO SRS solution.

We now have two main [inaudible]. The main one is in Dortmund but there's a failsafe in Amsterdam. [inaudible] not provided [inaudible] by CORE.

We also have several secondary servers and we try to offer variety to increase the availability and resilience [inaudible] possible attacks. We have one in Barcelona, another in Madrid thanks to [inaudible], three more in Europe, and two in the rest of the world. Four of them are using Anycast to increase that security.

The privacy. We tried to incorporate the privacy by default since the very beginning and our privacy model has five main characteristics. It has to be free for everybody, both registrars and registrants. It is the domain based. That means that if that domain decides to hide the contact letter, [it is all for contacts] associated to that domain. It has to be compatible. And the selection is done by the registrant at the register level. There is no direct selection with the registry. The main one is that we let the registrant choose what he decides to do with their privacy.

To be able to improve [inaudible] the model, we divided our registrants between legal entities, in government, associations, foundations, and business of all types and natural persons. The legal entities are forced to have their contacts disclosed, while the natural persons are divided between depending on the purposes of their domain.

The commercial ones have the same behavior than the legal entities, are forced to show their contact data. The non-commercial ones are the only who can choose between show or hide their contact data.

To support the privacy model, we had to do some work, prepare some legal framework. This was work for natural. And we did some consultation to the Catalan and Spanish data protection agencies. We also worked with the Article 29 protection party. We added the GAC recommendations.

There was [inaudible] an amendment to our registry agreement during the fall of 2012 and it was included [directly] in the new registry agreement on 2015.

Practically, that means that we, in collaboration with [inaudible] needed to create a new EPP extension which involves changes in the domain create, domain update, and domain name [inaudible]. We also had to change the output of our WHOIS.

We also asked our registrars to prepare their interfaces to offer to the registrants the right to decide about their privacy. It means a new domain has to define if it will have the contacts disclosed or not since the very beginning, since the creation. But, already existing domains have the right to choose if we need to change their privacy to hide, or maybe to show, the contact information. Our registrants had [inaudible] also. A big difference in the transfer requirement that I will explain later.

This is the formal syntax of the data extension. The most important part is we added a new type with two possible elements on the three operations and defining the registrant data domain as natural or legal. And in case of natural, we also incorporated an attribute to define if they decide to show or hide the contact information.

This is an example of how it was added the extension to [inaudible] operation. For example, [inaudible] dot-cat decided not to show their contact information.

Here you have [inaudible] modified our WHOIS output. You can see that the generic information is showed, but every level that implies sensible data is empty, is blank. We remain printing the levels without the fills concrete from that domain.

Moreover, we also added three possible contact forms to answer possible IP protection, law enforcement, technical, or contact between undisclosed registrants to cover all the possible questions or problems related to that implementation.

As I said before, the transfer of a domain, of an undisclosed domain, implies adding the [inaudible] information also in the domain name for [inaudible]. If the registrant doesn't add this information, they won't be able to get the contact handles, so they won't be able to contact the registrant in the actual domain as it is required by ICANN.

Moving to the adoptions, we get the first version available for testing during the fall 2012 and it was presented to the general public on

January 2013. At the time, we had the seven most important dotCAT registrars that was able to work with that extension.

After one year and a half later, more than 80% of domains had the right to decide about their privacy and now it is complete. Every dotCAT registrar has this right. Of course, the newly accredited registrars implemented it by default, practically.

Now we have around 54% of our domains with the legal privacy. The registrants are legal. Another 46% of natural individuals and that 46% are divided between 60% of domains without commercial purposes and another 40% with commercial purposes.

In consequence, we have around 72% of our domains disclosed and the other 28% undisclosed. During the adoption and [these years] of usage, we had some important issues. The main one is that the registrant still don't know this feature. We tried to explain them, mostly in the beginning, with an general e-mail to all our registrants, but after that, we hope that the registrants will do that work. [inaudible] accomplish.

Also, registrants sometimes doesn't have a direct option to define their privacy. The registrant has to open a specific ticket asking to change that setting. Moreover, many registrants are still asking to disable the privacy during the transfers. The registrars are typically recently accredited to the dotCAT market but also very large registrars which [inaudible] annual results related to this recently accredited registrars. Sometimes they set default policies. I mean, for example, each new create will be directly attached to a legal policy, for example, without letting the registrant decide during the registration.

Finally, in [two] concrete cases, this features were lost by the registrar during deploying of the new version of their site. Then we asked them about what happened and it as already incorporated again. But this is a concrete case that it good to know.

Anyway, instead of GDPR or many changes of this year, we decided to continue offering this option to our registrants. So, if it's still working and natural can [add] some information at this point. Please?

NACHO AMADOZ: So, what we're doing in our consideration is more respectful to the protection of the privacy of the registrants according to our interpretation of the previous [inaudible] on the current framework. We also agreed with the Spanish data protection agency and the Catalan data protection agency that this would be what we would be offering and we went through a very, very, very long process in ICANN through an RSAP procedure to get this approved by the board and the board confirmed their approval after a reconsideration request submitted by the Intellectual Property Constituency.

So, we've gone through many, many different steps to ensure that the kind of model that we are offering is privacy proof.

At this current point of time, and given the necessity to renew the temporary specification every 90 days and the uncertainty that the EPDP process is at the moment, we are still staying with this procedure that we established back in 2012. Done.

PEP MASOLIVER: So, thank you for your attention. If you have any questions, feel free.

EBERHARD LISSE: I would like to ask an obvious question from the chair. Is it that some registrants are not really interested in this because they provide their own commercial privacy provision?

PEP MASOLIVER: Come again, Eberhard? Sorry.

EBERHARD LISSE: You say some registrars don't use this. Is this perhaps because they have their own privacy providers for which they charge money? In other words, if they used your system, they would lose money doing it. Could that be a cause?

PEP MASOLIVER: I think that's the case because we've been telling them that they should be offering this, and at the end, most of them do it, even though it is a complicated process for a client of their to reach the part of the website where this can be activated. And it's much easier for that customer to just pay the small fee that it is required to have this proxy service offered. Yeah.

EBERHARD LISSE: We'll take on question for the floor.

GAVIN BROWN: This is Gavin from CentralNic. Contact objects can have privacy flags set on them. I was wondering why you chose the approach to apply the privacy settings to the domain rather than individual contact objects.

PEP MASOLIVER: It was a hard discussion in the beginning because our implementation also had the privacy flags for the contents, but we decided that our object is the whole domain. If the domain has, for example, commercial purposes and it is for a legal entity, it's clear that the contents have to be disclosed. But, on the other side, if it is from a natural, the purpose applies to the whole domain, not only a concrete contact. It's also easy to manage for the registrant.

GAVIN BROWN: Okay, thank you.

EBERHARD LISSE: What this means, of course, is that for individuals who are disclosed, who are resident in the EU, they must then specifically say that this information can be disclosed anyway. But [before it] we cannot disclose European residents' private data as a default. So, that's basically GDPR comes on top of this, so if a contact is disclosed, some information may not be disclosed because it's also not necessary. Okay, thank you very much. Also quite an interesting approach. Please give him a good hand. [applause]

We still have got some seats in front of here, so if somebody still doesn't like to stand, please. We now have a presentation from Maciej Korczyński. He used to be in the Netherlands but has now probably received a promotion to Geneva. Go ahead.

MACIEJ KORCZYNSKI: Actually, to [inaudible], but close.

EBERHARD LISSE: Oh, so sorry.

MACIEJ KORCZYNSKI: No, worries.

EBERHARD LISSE: It's my cataract. I can't read so far.

MACIEJ KORCZYNSKI: So, hello, everyone. Thank you for the introduction. Today I will present the attack that we refer to as DNS zone poisoning, not to confuse with DNS cache poisoning, and its root problem which is that DNS dynamic updates, non-secure DNS dynamic updates protocol extension and also our almost two years effort in notifying affected parties. I will discuss it also in the context of GDPR. So, we might like it or we might not like it, but we all need to adjust. Also, our notification strategies.

So, first I will discuss the attacks against DNS resolution path. I will discuss the zone poisoning attack and its root problem. I will discuss the requirement specifics and the threat model of the DNS zone poisoning attack. I will also give you a demo of the attack itself.

Then, I will discuss the global measurements and the affected domains and DNS servers, and then finally, I will present our findings about notifications in the context of GDPR.

So, on the right side, you can— [audio cuts off]

[END OF TRANSCRIPTION]