ICANN75 | AGM – Tech Day (1 of 4)
Monday, September 19, 2022 – 10:30 to 12:00 KUL

KATHY SCHNITT: …or moderator of this session. If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor. Please state your name for the record and speak clearly a reasonable pace. Mute your microphone when you are done speaking.

This session includes automated real-time transcription. Please note that this transcript is not official or authoritative. To view the real-time transcription, click on the closed caption button in the Zoom toolbar.

To ensure transparency of participation in ICANN's multistakeholder model, we ask that you sign in to Zoom sessions using your full name. For example, a first name and last name or surname. You may be removed from the session if you do not sign in using your full name.

With that, I'm happy to turn the floor over to Dr. Lisse.

EBERHARD LISSE: Thank you very much. So we do this now properly. I have introduced myself, so we don't have to repeat this. We actually do. We need to record this. My name is Eberhard Lisse. I'm the chair of the technical working group and the ccTLD manager for .NA. I haven't counted how many Tech Days we have done so far. So far we're approaching 50. Once we've reached that, I will obviously mark it in a nice color so that we know about it when we have it on the agenda.

First today, I'm going to go through the presentations a little bit so that we know what's going on and in what order. First we have quite a presentation .TW about DNS abuse and the Internet of Things. The speaker [Catherine] Ho is apparently now online. She had a little bit of an audio issue, so she will start.

Then Peter Thomassen who is sitting next to me will speak about automatic DNSSEC bootstrapping. I saw this on an IETF mailing list and I found it interesting. And since we speak the same language and I found out he was coming anyway, I invited him to give this presentation. I'm quite sure it's quite good.

Then Michael Bauland and Marc Blanchet are going to speak about universal acceptance for recognition systems.

The exact titles of the presentations may differ slightly. I tried to get them on line on the agenda. The actual presentation when

you download them from the website where they will be published may have slightly different titles.

The agenda will also be published on the website, and all the speakers in blue will have their mailing links clickable. So if you want to contact them, it's easy to do.

And then we have Fred Baker from ISC who will speak about query name privacy. That means what happens between two name servers when you query a name. He sits in the Pacific time zone so it's very early in the morning for him, but he indicated he is around. Next slide, please.

Then Ted Bartles from the ICANN technical staff or travel staff or meeting staff will talk a little bit how to organize these big meetings in the times of COVID. We have had a presentation about how they do meetings and how much stuff, how many containers they have and so on, and how they wire us up a few years ago. I think it was in Singapore. But now that COVID produces additional complications, I felt it was interesting to hear from them again.

Then Brett Carr who is a new member of the technical working group is also a member of the Customer Standing Effectiveness Review Team and will give us a brief report about that.

And then we have remote participation from ICANN technical engagement. Can I have the next slide, please?

Then we have again a little bit about registry best practices. I understand recently there was another African top-level ccTLD that got offline for a day or two due to issues with securing the DNSSEC keys and their virtual machines. So it's always good to hear again and again and again what best practices are in this regard.

I will try to ask my colleagues from that ccTLD to present on a future Tech Day because we are a friendly audience and we can all learn from these things. And we should all learn from these things so that we don't make these mistakes that have been made or these issues have come up that they don't come up again.

Then we always know about DNS abuse, but I found it quite nice to hear about a definition, what it actually is.

Then ICANN has a new program, KINDNS. It's basically best practices in the same line of best practices that we hear on the first presentation. So Adiel Akplogan will speak to us. And then the next slide, please.

Then we had one presenter abscond and one presenter sort of got bumped off his plane. So we had to add some presentations on short notice. And Eduardo Alvarez asked a while back about his EAI survey tool.

We are not so much interested in EAI. We have had lots of presentations about it. But we will of course hear a little bit of the results. I'm more interested in the tool. Survey tools are always interesting for us for our clients to gauge or judge their satisfaction, so it's interesting to hear.

And then probably our hallmark presentation will be Kim Davies. The IANA function operator is introducing a new root zone management system, and so they have half an hour to explain to us what's upcoming.

And then Brett Carr as the new member has been volunteered to give the usual closing remarks. We then sort of a tradition have somebody other than the chair review the day and say what he or she liked and not and so on.

There are a number of breaks. I am not too strict on the timing. So if we run into a break a little bit, it's not a problem as far as I'm concerned. But the breaks are mandated by ICANN, and I don't want to interfere too much with it. It's just to let you know if we run a few minutes over, I'm not going to be too difficult about it.

That said, can we have the first slide for the first presentation? There you are. I can see and hear you.

CHIA-LING HO: Okay, thank you, Dr. Lisse. Good morning to all Tech Day participants and greetings from Taiwan. It's my honor to be here sharing our team's work with you experts and interested parties of DNS. My presentation topic is to demonstrate some of our findings in our latest case study on DNS abuse and IoT devices and technologies in Taiwan.

I'm a policy researcher from TTC, Taiwan Telecom Technology Center, hence this presentation will more likely more policy review points rather than pure DNS technology discussions.

Our center's primary mission is to facilitate our government and private sectors in applying cybersecurity regulations receiving our cutting-edge penetrating tests and so on. Yet not to be fooled by our name, our work is not limited to the telecom industry but also other areas such as IoT devices industry which will be elaborated on during my presentation.

We also have a research division in our center where our work to collect the data and information from [inaudible] [mission] to our center. Therefore, collecting with national and international [economy] and research partners to provide policy recommendations for our ministry of digital affairs is also part of our work. Can I have our next slide, please? Thank you.

To begin with, this presentation I would like go through the contents of it. firstly, I will define our research title DNS Abuse and

the IoT, especially what part of DNS we're going to work on. And explain our IoT cases by narrowing down to those with critical infrastructure concerning transportation, such as IoT implementation [on] the airport, road traffic system, and [inaudible].

Secondly, our research objective is to respond to and echo the document of ICANN SAC105 on IoT and its threats to DNS along with introducing our research methodology, our cases, and our research findings. Next, please. Thank you.

The Internet of Things (IoT) is emerging Internet application that extends network connectivity and computer capabilities to objects, devices, sensors, and items not ordinarily considered to be computers. The IoT is expected to connect 20 to 30 billion of such objects to the Internet in the next decade with some analyses even predicting 100 billion connected objects by 2025.

With a topic scale like this, we cannot define the research area firstly to narrow it down. Since the DNS abuse definition is still pending and without any public consensus, we've taken a deliberate decision on our own on DNS abuse as anything that threatens the stability of the Internet and the DNS structure. Mainly cyberthreats on the [resiliency] of DNS.

As for the IoT cases that we choose for our study, we would like to focus on the most innovative IoT technologies applied to our

national critical transportation infrastructures, such as smart V2V, smart airports, and smart harbors [inaudible].

The SAC105 report of ICANN mentioned that contemporary IoT devices with an IP stack typically exchange data with one or more remote services hosted on the Internet and locate these Internet services using the DNS. We will be monitoring these while implementing our cybersecurity team's [excellent] work on labeling programs and certifications of DNS security. Next, please.

In addition, we are firmly aware that Taiwan has taken an important role in the IoT industry in terms of manufacturing and exporting IoT devices to the world. According to the Taiwan Vulnerability Note from Taiwan's Computer Emergency Response Team, we can have a rough image of how these devices can be part of our concern. Next, please.

Our data are from TWCERT a top IoT devices security response institute run by TWNIC, our domain name registry, allowing us to have a clear picture of our IoT device implementation risks relating to DNS threats. As you can see it includes a variety of devices and the list can go on and on. Next, please. Next, please.

So we are trying to show a map of both government and IoT technology developers on how to minimize these risks in the IoT

innovation process in terms of DNS abuse, no matter its [POC] concept, [POS], or [POB] services or business [stage].

Our objective of this presentation is to introduce Taiwan's current cybersecurity standards for IoT fields and establish a basic understanding of the level of DNS abuse awareness and decision-making policies amongst the multistakeholders of the IoT industry. And through our empirical studies to polish IoT field security regulations and standards. Next, please.

We are trying three different methods to collect our information. Firstly will be combining through the documents of IEEE and ICANN relating to IoT and DNS security [inaudible] with the latest international security standards, such as CMMC and NIST to generate our questionnaires and interviews factors to the targeted cases.

Our cybersecurity team is not only the first institution to gain the ISO 17025 and UL lab qualify or performing testing, sampling, or calibration but also focuses on IoT devices and technology security, analyzes [inaudible] standard [providing].

The goal is to help enhance the competitiveness of local developers and facilitate their products and services to meet international standards. Our research is to observe the adoption of these standards and guidelines while our cybersecurity examines the cases of IoT fields. Based on these steps, we try to

identify if there are any connections between IoT and DNS abuse, interviewing the responsible parties and monitoring their chosen solutions. Next, please. Next, please.

Due to our case study project was initiated by our ministry of transport, we have collected several IoT application cases focusing on critical transportation infrastructure such as vehicle-to-everything innovation on the roadside.

This case aims to maximize car flow capacity making traffic smoother. It's very extensive data collection and calculations. It's also innovative in facilitating emergency vehicles such as the first-response team to fire or car crashes, etc., to allow the traffic signs and lights to collaborate better to adapt to any urgent scenario. It combines both physical roadside units and cloud roadside units to process the data that the team gathered. It also uses a 5G router as a primary protocol gate to exchange data. Next, please.

In terms of smart airport we have a second category of our collective cases. So smart airport and this IoT innovation application. As we observe, two innovations have caught our attention which is the automatic car carrier and the auto [inaudible] [report]. With sensors on webcams the automatic vehicle is designed to pick up and drop off passengers between

terminals. Yes, we are all aware that some of the terminals are too big to go anywhere by foot.

Another innovation in the auto [report] that will sanitize the airport area by spraying the disinfect miss [carrying] sensors, cameras and infrared sensors as well. It can function only when no human is within the auto spray range. It is indeed innovation inspired by the aftermath of COVID-19, and it can make our world a safer place for traveler [hopefully]. Next, please.

Smart harbors, using 5G networks through base station on the ground and on the satellite orbit combining the technology of Wi-Fi 6 the drones with sensors to predict the courses of boats aiming to prevent them from collision by sending earlier alarms. In addition, drones can also monitor air pollution and oil pollution, etc. Next, please.

The innovation [background] is based on the fact that the ship with considerable inertia which is [deadly deflected] that they cannot turn quickly so collision often occurs. In the aerial photography operation of mobile UAVs images recognized by AI predict ship movement trajectory, and the control center will send out warnings to avoid collision. Next, please.

So the drone carried more than one IoT device, as you can see, such as a webcam, optical displacement sensor, and also an infrared thermal sensor with a camera function. The map on the

left on the slide is to show how the drone connects with the Internet. In order to overcome the disadvantage of [inaudible] webcam connection they use a 5G private network and public network. Meaning when the drone is close to the shore by let's say 750 meters, it would change connecting method from Wi-Fi to 5G LTE. Next, please.

From the above [cases] we learned that the characteristics of a 5G network, low latency and broadband capable of massive machine linkage, are ideal for the innovation of IoT deployments. Yet they were emerging DNS abuse concerns while using the 5G technology. This map is to demonstrate that 5G complexity generates more risks. Next, please.

This map breaks the 5G IoT complexity into layers, including sensing, network, and application layers. Where relevant, hardware, software, and firmware technologies are deployed. However, the vertical application of IoT technology [avoided discuss] by the ICANN SEC105 document which explicitly expressed that the IoT and DNS report of ICANN means to narrow down into horizontal layers. From the complexity of it, the ICANN report [needed] to be reshaped and redirected quickly. Next, please.

Although the shift from 4G to 5G can be risky for the integrity of the Internet, there are still bright sides to be discovered. Using a

private 5G network can prevent DNS from exposing [under the] threats as an all-IP network like 4G. It's a great way to limit threats from the public network if proper actions have been taken. Using a private 5G network can be one of the solutions for dismissing DNS abuse and threats. Next, please.

Our national cybersecurity strategy has been shifted these days as well. It's our great concern too in terms of cybersecurity from the national security point of view. Under the circumstances and stress of being constantly attacked by other Internet stakeholders we have no intention of letting this big chunk of the industry at risk. Next, please.

In terms of our center, our center has helped draft—next, please. Yeah, thank you. Helped draft the guidelines for security assessment. Also, it has a certificate testbed for accrediting IoT technology developers to certifications, including [examining] safe DNS measures when necessary. Next, please. Thank you.

We do have our own systemic marking program for both IoT devices and the IoT field. IoT devices can be seen as a separate component of the whole IoT field application. Yet each device has its certification. Meaning in order to gain a certificate each one of them would have a specific set of criteria to meet. Next, please. Thank you.

The first end-to-end security assessment guideline for 5G IoT applications covers the sensing, transport, and application layers. First, we must [build out] threat models and proceed with vulnerability and penetrating tests. The outcome will be used to draft an impact analysis, and a security threat report will be generated. Next, please.

So in detail of this guideline, from Chapter 5 to Chapter 8 the guideline lists all the detailed elements to be considered by the engineers working in our cybersecurity team. There are 20 threat models, 46 security controls, and 10 vulnerability tests along with 19 penetration scenarios and 18 penetration tests to be run. Finally, there are 2 impact analysis models to comply with. Completing the test and gaining the Level 3 certification will take more than six months considering the quantity of assessment criteria. Next, please.

We have to emphasize that as of today Taiwan's security guidelines and certification for the IoT field is a pioneer innovation relating to IoT security. The only one that exists in the world. Yet for the moment there are no regulations that strictly enforce IoT devices, applications, and developers to follow. We're working on a more inclusive approach to encourage the private sectors and communities to invest in cybersecurity willingly.

We can see from the comparison table between device and field certification that the latter was a more inclusive capacity to have a complete [inaudible] cybersecurity examination. IoT devices exist for, let's say, hundreds and thousands of different yet the security certification will never succeed in covering them all. So we are using IoT field certification as a resolution to cover all aspects in the application field on the IoT deployment area. As you can see the pros and cons of two different certification systems in the table. Next, please.

With an increasing worldwide [trade] security requirements, our centers security [lab] targets to be accredited by the international security standard organizations and assist Taiwanese manufacturers to comply with the security requirements, such as U.S. standard, U.S. lab, and ISO 15408 common criteria for information technology security evaluation CC testbed as well. Next, please.

EBERHARD LISSE: Can you come slowly to an end of the presentation?

CHIA-LING HO: Okay.

EBERHARD LISSE: I'm not going to rush you too much.

CHIA-LING HO: Okay. According to our research some of the in terms of industry stakeholders are taking DNS threats seriously and they urge the government to put them into the cybersecurity assessment requirements. Yet the [impasse] of this topic remains. It's about the nature of the industry of Internet of Things which lacks resources. We know that the IoT industry comprises three providers: hardware, software, and firmware. Unfortunately, except for hardware equipment manufacturing neither software nor firmware service providers have enough market share to influence the related policies.

On the other hand the price of IoT devices has significantly decreased. Therefore, to be more competitive in terms of market price they usually will not invest in security let alone DNS abuse prevention. According to the recent report by Nozomi Networks labs in May this year another example why IoT devices are a risk for DNS. A flaw in all versions of the famous C standard as everybody should know can allow for DNS poisoning attacks against target devices. The scope of the flow is vast as major vendors [inaudible] distribution use these devices.

Meanwhile, it's designed to open WRT, a common OS for routers often deployed throughout various critical infrastructure sectors. So IoT application developers often have limited resources,

especially those in Taiwan. Therefore, they would face pretty much the same problem by using open resources from open libraries.

Once the coding is tested to be [inaudible] they don't even have the abilities to amend it. As for the [botnet] targets to DNS the device manufacturers also have the problem as IoT developers. They focus more on moving product and boxes and selling them rather than making devices safer.

According to our interviews some say a domain name is like a basic infrastructure that the national government own institution provides. Security relating to a domain name is not as important as other factors.

So the measure to be taken by the IoT developers depends on the level of the IoT testing field. When it comes to [POC] concept, many new tech creators are all about ease to make it work technically speaking. So security is the last thing that comes into their mind. Yet when it comes to POS, product of [inaudible] services, or POB, product of business, it would depend on factors such as investment, how important this innovation will take place, and who their end users are.

There is no difference between other IT products and IoT products market whether the developers willingness to take security seriously depends on the [inaudible] elements. So IoT in

the 5G era is becoming a challenge for cybersecurity, both in its scale and depth. According to the [inaudible] research and other prestigious research findings, we can take a closer look at the structure and place DNS security on the [grand] map of IoT cybersecurity threats.

In this presentation we tried to merge both affects of 5G and IoT [utilization] on the DNS abuse issue and present the efforts of our center and our team's work on minimizing the risk of undermining the stability of the DNS structure.

So about the challenges, the according to the ICANN report SAC105 the IoT is an opportunity because IoT devices send and act upon physical environments and will therefore have new security, stability, and transparency requirements that the DNS functions can partly [inaudible].

Personally I think DNSSEC is one of the best ways to solve this problem. Nevertheless, the [inaudible] sectors feel about the about the DNS is not too optimistic in Taiwan due to the disadvantage of the mechanism of being in the world of capitalism. That is to say when there is profit to gain in terms of DNS security, thousands of different products emerge and DNSSEC would not be able to deploy fully.

During the GNSO policy update of the prep session last week the chair of the GNSO Council, Mr. Philippe Fouquart said, and I

quote, “The topic of domain name in the IoT has been debated somehow during some of the plenary sessions that we had. Yet there has been no specific policy work under the GNSO on this particular topic on how DNS can be used for IoT. There is certainly no current path I’m aware of on this particular topic.” So now we understand why it’s at an impasse in finding common ground on the solution.

EBERHARD LISSE: So you must really come to a close now.

CHIA-LING HO: Yes. The [inaudible] for that IEEE and ICANN published on IoT devices and DNS abuse have several types of connection. In Taiwan we’ve mentioned the [inaudible] devices and their risks on implementation abroad without taking proper measures to reduce the risk of DNS threat.

Secondly, the lack of awareness among the IoT technology end users and developers that the primary concern is [inaudible] the cybersecurity of the products but still with the [inaudible] of profit.

So it’s up to our government and top industry management class to consider cybersecurity when it comes to innovation

technology such as 5G and IoT that make cybersecurity a physical safety scenario.

We also strongly recommend to all cybersecurity center contributors experts [inaudible] and institutions and governments to add DNS security examination into their authentication process and certifications. It's the best way to consider the reality of our capitalism [inaudible] society and the market value of the security business. Next, please.

This is the end of my presentation. Thank you for listening. Sorry for the extension.

EBERHARD LISSE: Thank you very much. Not a big problem. I just want to keep a bit of the agenda flowing. Thank you very much. Interesting topic. A little bit over my personal head, to be honest. But always good to hear current developments. I'm not going to allow any questions at the moment. If we have questions at the end, we can maybe make a plan. Peter Thomassen, you have the floor.

PETER THOMASSEN: Yes, let me quickly share my screen. All right, hello. Good morning. I've been asked to give a presentation here about DNSSEC bootstrapping, so how to turn on DNSSEC for a registration that so far does not have DNSSEC yet. And there are

various ways of doing it, but none of them so far was [ultimately] authenticated.

The topic has been talked about in the IETF and in some specialized DNSSEC sessions at ICANN meetings. And it's in production now in some places and probably will soon turn into an RFC. So it's time to make it accessible to a bit of a broader audience, which is why I'm here.

You may know that the DNSSEC validation rate according to APNIC is about one-third. So chances are that your ISP or your resolver does validate DNSSEC signatures on responses when the signatures are present. But unfortunately, often the signatures are not present or the chain of trust is not extending down to the domain that you are querying because only 6% of delegations are secure. The sources for this are below.

Of course, it depends on the region a bit. Sweden, for example, has both very high validation and also secure delegation rate. But generally, there is a disparity between the validation rate which is now significant, about one-third, and the secure delegation rate which you could say is still close to insignificant.

At deSEC, we do manage DNSSEC hosting and we sign all our zones by default. And we notice that even if we do that, registrants often don't put the DS records into the parent zone. So even if you sign, you only get less than 50% of the delegations

actually secured. So that is something that doesn't sound like it should stay like that forever.

The question is, why are there so few delegations that are secure? There are multiple reasons probably, but one big reason is that deploying DS records in the parent domain is multiparty problem. The DS record has information about DNSSEC keys that are used in the child domain, and those keys are maintained by the DNSSEC signer who is usually the DNS operator. So that's the origin of this kind of information, and it needs to go into the parent zone. So you somehow have to get it to the registry.

But usually, the DNSSEC signer, so far at least, doesn't talk directly to the registry and it often involves the registrar as a messenger and/or the registrant who has to go to some web interface and ask the DNS operator for cryptographic parameters they don't understand because registrants are usually not technical people then put things into weird forms. And there's different formats and dropdowns and whatever and it's very complicated, so nobody does it essentially. At least nobody does it when they're in their right mind or nerds like us.

So it's error-prone. It involves too many parties. It's usually slow because it's manual. It's often out of band which in itself is not strictly a problem but it adds dependencies to other layers and complexity. And it's not properly authenticated, at least usually.

There are a lot of parents, for example, that just will do unauthenticated CDS scanning and stuff like that. So it needs automation, and ideally it should resolve all of these problems.

Also, any automation must involve the source of truth which is usually the DNS operator. Because if you don't have that, if the automation only reaches between the registry and registry for example, you don't gain very much. So we need some automation that allows for independent participation of the DNS operators.

Here is an illustration of the traditional deployment of DS records for the first time when you want to secure your delegation. So what happens is that the DNS provider at the bottom here, the service provider who is usually also the signer, will do whatever magic. They will be signing the zone and put that into their authoritative servers.

And then it gets complicated. The registrant goes to the DNS operator and asks for these records. And then if they obtain the domain name through a reseller, not directly to the registrar, then the reseller is involved also. Perhaps they don't support it. And hopefully you don't have a reseller. You can talk directly to your registrar. Blah, blah, blah. Very complicated.

And eventually through EPP and registry and whatever it ends up at the TLD server. And then you have this link between the TLD

server and the data that's on the authoritative server of the child. So that's all very complicated.

And if you look at the structure of this, there is the DNS hierarchy levels which are the child at the bottom and the parent at the top, and those are all technical parties in a sense. And in between, there is the registrant who the whole thing is hinging on the registrant. And the registrant isn't even technically part of this hierarchy, right? which is why the field on the left is empty. There is not really a place for the registrant except that they're part of the deal.

So it would be cool if we could just get rid of this because otherwise it's too complicated. And we actually know that and it's time to change it. So how can we change it? We need to somehow establish trust first to the DNS operator. Let's see how we can do that. And then once we have that, we can transfer that trust onto the child domain. Let's see how it can be done.

The Internet draft is linked at the bottom. I also have it in the last slide. It's also in the title slide if you download the slides.

So how does it work? We need three things. First, we need a way for DNS operators to publish the DS information about the zones they are managing in a secure way. So we need a publication signaling kind of mechanism, and it needs to be on a per zone basis. For each of their customer zones, they need to be able to

announce independent things in case they use different—yeah, whatever. The DS records are just not the same even if you use the same key for a different [domain], so it needs to be per zone.

We do that using a namespace under each name server host name. So let's say, for example, NS1.deSEC.io in our case, we can use subdomains of that and put the customer domain there. So it could be example.com.NS1.deSEC.io. And then we will also add intermediate labels for logistical reasons, but essentially you put the customer domain in front of the name server host name.

And under that name you can publish whatever information you want. As the DNS operator, it's under the operator's control because it's in their name server zone. And you can require DNSSEC for DNS operators that implement this on their zones that have their DNS name server host names. So we, for example, use DNSSEC also for NS1.deSEC.io. And because of that, we can publish that stuff under subdomains of that.

So that's the general signaling mechanism. It can be used for all kinds of things that you want to announce about to your customers. And so far, the only reasonable usage I know of this is to ask DNS operators to publish authentication signals for CDS and CDNSKEY records.

So CDS and CDNSKEY records are usually stored in the child next to the SOA record. It is where the child has a preliminary version

of the DS records they want to deploy in the parent. So the C stands for child. The parent can look at that and fetch it and then deploy it in the parent.

The problem is so far there was no authentication for this when you were doing the first deployment of DS records because, well, it's the first. So you have a chicken-egg problem. For key rollovers you can use the existing chain of trust for the child, but for the first time it's difficult.

So we propose to use the signaling mechanism from the bullet point above and use that to publish a copy of the child's CDS records under that signaling name under the host name of the name server. That will be signed with the name server zone's keys. And then the parent can come and fetch that and validate it. And if the CDS records in the child and the ones announced by the DNS operator agree, then they're endorsed by the operator, cryptographically verified, and you can go ahead and provision the DS records at the parent. So you transfer the trust that was preexisting from the DNS operator to the child.

Now this was a lot of words and a lot of black on white. Some people like green on white so there's an illustration. Let's say we want to DNSSEC for example.com using the provider provider.net. Let's say the provider already has DNSSEC for their own domain including their name server domain,

NS1.provider.net. And we will also prepend this with an extra label that's called _signal which is the roof under which we run the signaling mechanism that I mentioned earlier.

Now somebody does a registration of example.com. It is not green because it doesn't have DNSSEC yet. And NS1.provider.net is their DNS operator. The DNS operator now puts the CDS records, which are the tentative DS records for the parent, into the child zone. And at the same time, they publish a copy of that in a subdomain of their name server host name zone. So we also have a signal type identifier in front in casae there will be other signals in the future. Who knows. So it would be _dsboot.example.example.com. and then the things that references the provider and already has the chain of trust.

Now the parent comes to the registry or registrar. They look for this stuff. They validate it against what's in the signaling zone. And ta-da, you're done. I don't know if you noticed, but due to the simplicity of that it might have illicited you. No, that's the wrong word. Whatever. You might have missed that the registrant actually didn't even appear, and that's the beauty of the whole thing.

This already existed as RFC 8078 without the left part. So without the authentication some TLDs already do it. I think more probably

could be convinced if it would be more secure with this mechanism.

So we use an established chain of trust on the left to take a detour and we find identically co-published there the CDS records from the child. They're authenticated, immediately available, and at least under the normal assumptions of DNSSEC you can't have an active on-wire attacker undermine this.

So cool, right? hopefully. So if you do that, you can get rid of this whole intermediate layer. And deploying DS records for the first time just involves the DNS operator putting stuff on their server which is the CDS records in the child zone and also under the name server subdomains and the registry fetching it or the rather, whatever. It depends on the setup. But the parent entity fetching it and then putting it on the TLD server. I think these steps are actually the minimum number of steps required. So I don't think it becomes much simpler than that in the future.

It's already in production. There are two DNS operators that I know of that have this for all DNSSEC-enabled domains already. One of them is deSEC and the other is Cloudflare who announced this at the last ICANN meeting in The Hague. They cover 23% of the top million domains according to the Tranco list. And if all of the parents of these domain names would also be able to do the scanning and do the authentication for the DS records, then that

could give quite a boost to the 6% of secure delegations that we had earlier.

But on the parent side there are only two ccTLDs which support it so far, which are actually one registry. So that is Switzerland and Liechtenstein run by SWITCH. Also, Chile is close to rolling it out but hasn't quite done so. So this is including the authentication mechanism, the lefthand side detour.

And there are also five other ccTLDs, Czech Republic for example and Sweden, who do the bootstrapping CDS scanning from the child but without the authentication. So that is good but not as good.

And then in the future, GoDaddy is planning. I mean, they are planning now to in the future introduce DNSSEC bootstrapping for the domains for which they are a registrant in case the registry wouldn't do it. So that together with Cloudflare probably would lead to quite some overlap between child and parent side and give some significant impact.

I see some chances that there are some TLD managers here, so you're invited. For example, I'm sitting next to one, I think. So everybody is invited to adopt this. The specification is underway to the last call in the DNSOP working group in the IETF. I don't expect any significant changes. There is a link here if you want to get into the specification. It's actually not very difficult.

Client-side extension implementations are deployed for a significant number of registrations, and we need parent side implementations. So for the ones that already do the scanning but without authentication it would be cool if the authentication would be added. But I understand it's like changing existing systems and all of that so it's complex. It might actually be easier for those who are starting the CDS scanning from scratch to start out together with the authentication.

There is some example code available for this that is thanks to RIPE NCC which also does CDS scanning for their reverse DNS stuff that was made available to me. It's in GitHub. And if you like and if you need something to get running to get up to speed as an example, you can approach me. Reach out to me. I'll send it to you and we can work on improving it or whatever, do a tutorial. It's all feasible to do.

So one we have these registrations DNSSEC will be easier for the whole community, and that's I think what should be the goal. Thank you. Any questions?

EBERHARD LISSE: Thank you very much. As you may or may not know we are on [inaudible]. So I will speak with the developers from [inaudible] and see what they say about it. Because if they implemented for us, then it's implemented for all the ccTLDs that use it. The

demand for DNSSEC in Namibia is zero. So only our own company domains and my private domain is signed. The banks firmly believe that HTTPS is all they need. We can talk to them 1,000 times. They don't care. So it's a long-term project to raise awareness. But if we offer the service, if it's easy to do, then it's probably easier.

Are there any questions from the floor? I will take one question from the floor. Please take your mask off then it's easier to understand, and identify yourself for the record.

LARS-JOHAN LIMAN: I'm not sure if the mask is the problem with my English. Hello, Lars-Johan Liman from Netnod. I just wondered exactly which piece of information leaves the parent to look for the signaling at that specific record.

PETER THOMASSEN: So your question is, what is the trigger for the parent to look?

LARS-JOHAN LIMAN: No, not the trigger. The trigger is that it's built into the software, but the software needs to find the _signal.NS1.provider.net. How does it know NS1.provider.net?

PETER THOMASSEN: Oh, because the parent has the NS record set from the delegation. It always knows where the child is delegated.

LARS-JOHAN LIMAN: So it uses the registry database information essentially for that?

PETER THOMASSEN: Yes.

LARS-JOHAN LIMAN: Rather than information from the child.

PETER THOMASSEN: Well, it uses its own registry database to contact the DNS operator.

LARS-JOHAN LIMAN: Okay, but it's the child data that's authoritative, right?

PETER THOMASSEN: Yes.

LARS-JOHAN LIMAN: Just checking. One last question. How does this make things less complicated?

PETER THOMASSEN: Well, you will see that it scales unlike the registrant manually….

LARS-JOHAN LIMAN: It's easier to automate. I will give you that. But it adds complexity to the system.

PETER THOMASSEN: But customer support is also complex.

LARS-JOHAN LIMAN: It is indeed. Thank you.

EBERHARD LISSE: I always say DNSSEC is easy, it's just complicated. Thank you very much. Nice presentation. Interesting topic. Kim, do we have anything from the remote side?

KIMBERLY CARLSON: We did have a question from Dirk, but it's already been answered. And that's all we have, so we're okay.

EBERHARD LISSE: Thank you very much. Thank you both our previous [inaudible].

UNIDENTIFIED MALE: I see your raised hand on Zoom.

EBERHARD LISSE: I beg your pardon?

UNIDENTIFIED MALE: I saw a raised hand from Warren Kumari on the Zoom screen.

EBERHARD LISSE: Zoom is in the…Warren I think was in the room, so he can talk locally. And leave this. Our staff is quite capable. Thank you very much.

Okay, Michael Bauland and Marc Blanchet, thank you very much for giving us this presentation, and you have the floor.

MARC BLANCHET: Good morning, good afternoon/evening, good night for people remote. It's great to see you in person. This is presentation about Universal Acceptance Roadmap for Domain Name Registry and Registrar Systems. I'll be presenting the first half, and Michael will present the second half. Next slide, please.

So we'll do a bit of an overview of the study, some analysis, gates. And then we'll discuss, Michael will actually discuss test cases and an example of registry and registrar test cases. Next slide.

This is a study that was done over the last six months, I guess, to help domain registries and registrars to make their systems support universal acceptance. What do I mean by systems is, for example, their registration systems which handle EPP, RDAP, Web, customer support, DNS zone generation, everything else.

When I say universal acceptance it's either or both support of IDNs, internationalized emails, long and new TLDs. Therefore, an example of this is if you receive an email from a customer and you are a registrar or a registry and the customer has an EAI address, does your system work? If the registrant email address is an EAI, does your system work?

The report is currently in ICANN public call for comments and the presentation URL is there. It's actually on the front page of the ICANN.org main page right now for the current public call for comments. And it closes on October 17, so I encourage you a lot to look at the report and place your comments. Next slide, please.

This presentation will essentially describe what the report is at some level. Obviously, there's more details in the report.

So the methodology of the study uses the UASG026 UA Readiness Framework which is a framework that was defined for a way to look at your systems, your software and identify what's the places where I should care about UA readiness.

That framework is generic to any application. So it's actually usable for a mobile application or a backend of a system for ecommerce or something. What we did for the study is actually apply it to a generic model of registry and registrar systems. You will see later on the generic model we used.

This model works for both gTLD and ccTLD. However, for example for gTLDs it actually includes a bit more details on the specifics for ICANN contracted parties requirements such as, for example, exports. So continuing on, exports for example, it says which fields you should be looking for that may need to be UA compliant.

That methodology then finds gates within the systems where UA support needs to be verified properly. It proposes some test cases for this verification. [Humbly], we could spend 100 pages on test cases. Michael will provide some examples. The report has more examples. But obviously, it's not completely comprehensive.

The good part is other UASG documents actually have specific reports on test cases for UA that you can look at and are referenced from this report. The report also analyzes two registry systems and one registrar system as examples. Obviously, the idea is not to pinpoint people but to actually help illustrate the process, the roadmap to make those systems UA ready.

And the registry systems we use Google Nomulus example because it's open source, it's more easy, and the KnippTANGO registry service and the registrar system COREhub, GatewayNG which both are developed by Michael and his team. The report is targeted to registry and registrar operators, registry backend providers, developers and technical managers. So it's a technical by nature. Next slide.

So, going a little bit deep into the report, will show you in the next slides registry high-level architecture, the registrar high-level architecture. And for each identified gate, the report describes the expected behavior of the software. We'll just illustrate a few of them for your for this presentation. Next slide.

So this is the registry high-level architecture. Obviously, it probably should look familiar to you, but it's a high-level abstract architecture, so adapt to your own environment. You know, there's no, obviously, considerations here in terms of cloud and stuff. It's just functional blocks.

And you could see in the picture that there's gates, G1, G2, 3, 4, 5, etc, that identify where the those gates are, the place where you would put—will do your test cases, and will identify what needs to be done.

For example, I'll take just one on the slide G8, is the interface to database. Well, we would obviously agree that you would make

sure that your database is UTF-8 or support Unicode because there are Unicode in those strings. Well, that's an obvious one, but you know, just to give you an idea.

So that's the registry one. So WHOIS, RDAP, the registrar admin interface, the registry admin interface, EPP interface, DNS zone generation, exports to third parties, email service and corporate web. Those are kind of a WHOIS. So, those are kinds of high-level functional blocks that we look into. Next slide.

This is the registrar one which is very similar, obviously, with some differences. The EPP interface is the client side, the registry side is the service side of EPP. But we use the same numbers for the same gates so that they're essentially the same thing. So in terms of UA compliance testing. Next slide.

So gates are numbered, unique for both architectures. Most are identical, but some are different. EPP usage is obviously different from registry and registrars. This is a generic architecture, adapt according to your environment. And gates were identified based on the readiness framework model. That is where the small picture is defined here. Next slide.

So for example, G 10, and G 15, identify exports to third parties such as ICANN. So there's a list of relevant fields and those exports that are identified and the expected format. G7 identifies the backend report, discusses important considerations about

backend development and the fact that, for example, some language libraries and open-source software may or may not be UA compliant, therefore affecting the backend as a whole. And we actually provide you some links to the different open-source libraries that are compliant and the level. There's another UASG report about that. Next slide.

We discuss about other considerations, protocols such as EPP, WHOIS, RDAP, generic considerations, string normalization, support of different scripts directionality, for example, IDN handling UTF-8 versus Punycode. Next slide.

That's the key here to remember if you look at this report. We're not talking yet about IDN invariants. being different IDN labels that are considered equivalent for registration. The reason here is that the whole policy, at least in ICANN environment, is not yet fully done. Well, the technical part will probably not change, but to be 100% sure. Obviously, the impact of variants on your systems is pretty significant, because your data model everywhere is very significant change.

So, we have put some considerations for you if you start tackling that problem, but it's not fully studied and discussed. So obviously, hopefully, in the next versions of that report, we'll consider the whole discussion on variants. Next slide.

Next section will be handled by Michael. And again, the whole report is including appendices that are presented by Michael is on public comment. Comments are due by October 17. Please read and provide comments.

MICHAEL BAULAND: Okay, thanks. Next slide, please. So, I will tell you about some more details of the actual test cases we did. I'd start with the choice and selection of tables. Next slide, please.

We, for testing, set up a sample registry system, serving the TLDs dot example and Japanese version of that string, I can't pronounce it. And for the labels we used to build domain names and email addresses, we actually obtained these labels from the IANA root zone database of example strings in different scripts. For example, the three shown here, but we use also other ones. And to be able to actually test email sending and receiving, we of course, needed to use actually existing domain names in the root zone and build out from that. We use the email address Michael at some Arabic string, which I'm told is something like ICANN dot Bharat. And another one we use is grüün@knipp.de and for those who know German, that typo was on purpose. Next slide please.

So for the registry test cases, we tested the EPP interface, the control panel, the web interface, DNS name server, Port 43 WHOIS, RDAP and escrow export. And next slide please. And for

this presentation, I'll show you one example test case, namely the EPP contact update. Next slide please.

For that example test case, we try to update or contact and set its email address to the one shown there, built from those labels I mentioned before. And it was a bit surprising that our Tango registry system did not allow that update, we received a parameter value policy error. Next slide please. So, we altered that test case and used an all Chinese email address and that one was successful. Next slide please.

So the analysis of that problem, we tried to find out what actually went wrong. So it was not that only ASCII characters were allowed as can be seen with Chinese email address. And we also found out it was not the Tamil script as a TLD label as such, because when we altered the email address by just removing some of the characters in the TLD, it suddenly succeeded. And we then debugged the code. And for the email validation third-party library javax.mail was used and that one simply marked the abovementioned address as invalid. So, the question is, what can we do about it? And next slide please.

It turned out that the library was working when we did not use the U label version of the domain name but the A little label version, so the same email address within an A label as a domain name was successfully valid data.

So the solution was to still stick with javax.mail validation, but implement a small workaround in that we first validated the domain name part of that email address individually using a different library that also works correctly with some Unicode strings. And if that domain name was valid, we convert it into the A label version, and then feed that email address into the javax.mail library. And if that email address is then validated successfully, we store the original email address. So using these additional steps, now Tango successfully validates all email addresses as far as we tested. Next slide please.

For the registrar test case, we tested the CORE GatewayNG system. Next slide please. The test cases are quite similar. Instead of EPP, we tested the proprietary API, core provisioning protocol, and again control panel, DNS name server, port 43, RDAP and escrow. And here we also tested the actual email sending and receiving using the WHOIS accuracy protocol and transfer notifications. Next slide please.

As a sample test case in the registrar system, I will show you contact create example. Next slide please. And we again started with the same email address that failed in the registry system and it was not a big surprise that this email address also was not successfully validated in the registrar system as both systems use similar code bases. So, that problem was expected. Next slide please.

But then, we were up to a surprise that the Chinese email address which was valid in the registry EPP request was also invalid in the control panel of the registrar system GatewayNG. Next slide please. And we then tested Latin script email address, Michael.mag@grün and that was also not successful. So this was a surprise to us. Next slide please.

For the analysis, we then realized that presumably any non-ASCII character was rejected in the web interface. Further analysis showed that the validating library as part of the vue.js framework uses a rather complex regular expression. But even though it's complex, it still only seems to accept ASCII characters. Next slide, please.

As fix for that problem, we thought about fixing the regular expression, but very quickly thought that this is not the best way to do it, it's far too much work to implement it in a regular expression. And it's far too error prone. So the next approach was to just use the back-end Java code to do the validation. Because anyway, every request coming in via the REST interface has to be validated, again, by the backend, to ensure that we really just store and work with valid data. And for that, we simplify the frontend validation by just checking for a very basic error, i.e. whether we actually have an email address in a syntactic way. So we now just check whether the string is some string at some string, dot some string with no real restriction for the some string.

And the final, more detailed validation is then afterwards done in the Java code, which then successfully validates the email address, similar to the example in the registry case.

Yeah, and that's our quick example of the test cases. All test cases can be found in the document which Mark linked in the first part of the presentation and which is currently in public comment phase, right? And feel free to comment and ask questions. Next slide. That's it. Thank you.

EBERHARD LISSE: Thank you, thank you very much. Any questions from the floor? We can take at least one. That is not the case. Thank you very much. Let's give him a good hand. Next presenter is Fred Baker. I've seen him shortly on Zoom. I can see him on Zoom. So he's waiting for us. The slides are up, you can start. Thank you, you have the floor.

FRED BAKER: Okay. Thank you. What I wanted to talk about is DNS privacy. There are several different approaches to that, that have been discussed in the industry, DNS on TLS, DNS on HTTPS and query name minimization. So this, like the other two, is now published as an RFC, and you can go read about it as much as you like. Next slide, please.

So I want to point it out that DoT and DoH are getting a lot of discussion these days for DNS privacy, but we don't see a whole lot of people talking about QNAME minimization, and we think that's unfortunate.

So, what do they do? Well, Dot and DoH are each initially targeted from a stub resolver to a recursive resolver, which gains in privacy across the backbone. But it means that certain systems have to trust each other. There's some overheads involved, such as [GCP.] Key Management is important. And you have to kind of know who you're talking to in terms of configuration. Next slide, please.

So TLS encryption is essentially using TCP with TLS in order to run DNS on TCP, which is, we've had that for a number of years now. And it does have some issues. One of them is the need for configuration and the amount of memory that is kept around for the different servers' session status. And there's an attack, which is to know the key in use. If you can identify the key in one way or another, you may as well not encrypt it.

As a DNS operator, I'm also concerned about bandwidth. We don't know how much bandwidth would be used. And it'd be interesting to know that. And then, of course, encryption and decryption imply a certain amount of computational overhead as well. So next slide, please.

Using DNS over HTTPS, which is what's used in Google and a variety of different services around the place. Once again, the clients need to know what servers are available for it to use, which is a configuration thing. And each server has to maintain session state, which is essentially a memory thing. And we have the same attack. If there's a way to know the key in use, all bets are off.

With bandwidth, HTTPS sessions use more overhead bandwidth and memory than typical UDP sessions. So just the fact of switching to HTTPS increases the bandwidth in use, and increases the amount of session computation. So next slide, please.

The question that we find ourselves asking is, what if we don't send the data? If we don't send the data, then there's no question of knowing the key and use some accidentally figuring out how things might work. Query name minimization could be used with any transport protocol, traditional DNS, DoT, DoH, DoQ, whatever. It works from the client or stub resolver all the way to the authoritative name server if you want to take it that far. It does require parsing of the query name by the servers and resolvers on the path.

It also uses a specialized port number which means that it's identifiable. There's the question of the cache organization. If you are keeping your cache sorted by the name that was passed along

and therefore would be returned to you, then you have to change the software in order to do that.

The attack, if any, is different, you have to have an intercept point which is before the label is removed, which is a little bit different. Next slide please.

Now providing privacy between recursive and authoritative servers is good. Privacy is good whenever it occurs pretty much. However, there are some infrastructure upgrades that have to happen for that. Next slide.

Now when I talk about a DNS query string, www.example.com would be an example of a query string, what QNAME minimization does is that the recursive resolver only sends the number of labels necessary to recurse through the hierarchy. So it might send that com to the root, it might send example.com to .com, it might send www.example.com to example.com. But it's not going to send information in other directions that could be observed and people could do nasty things with it.

So observers on path can view the query up until the point where it gets edited. So we're minimizing that information as much as we can. The authoritatives only see the part of the query that they need to process the request. As a root operator, what that means is that we're probably going to see less requests because they

would get picked up by caches earlier. And that's presumably true of other servers as well. So next slide.

So where we do deploy this, well, you'd only do it on the recursive resolver. You might have several recursive resolvers on a path. But that's where you would do it. If you want to use DoT or DoH, you can combine it with that from stub to recursive, it can be combined going from the recursive to the authoritative but there's no end user requirement. And so, clearly, there is some cost to anything that you do and QNAME minimization is not different in that regard. It does introduce a small cost, but it doesn't introduce the bandwidth or parsing overheads, encrypt-decrypt overheads that the other alternatives use for signing. So next slide, please.

So our recommendation look at this, we're actually very interested in QNAME minimization. We suggest configuring recursive resolvers with QNAME minimization, if it's at all possible. And users of course can select trusted recursive resolvers that they want. And the authoritative servers can monitor the effects of QNAME minimization. So that's kind of what we would hope that people might use this for. And I believe that's my final slide. So next slide, if there is one. Okay, I got that. Any questions?

EBERHARD LISSE: Thank you very much. Are there any questions from the floor? There is one from the floor.

BARRY LEIBA: Hi, Fred. Thanks. You talked about QNAME minimization imposing a small cost on the recursive resolver. What is that cost?

FRED BAKER: Well, that cost is the overhead of doing the change to the record as it's being passed.

BARRY LEIBA: Okay, so really, really small cost.

FRED BAKER: Yeah, very small.

BARRY LEIBA: Okay. Thanks.

EBERHARD LISSE: Brett Carr.

BRETT CARR: One of your slides mentioned that QNAME minimization required a specific port, but that's not how I understand it. I thought it just went over port 53 as normal.

FRED BAKER: Okay, I could be wrong on that. I believe that it uses a special port. But I defer to the document. Okay, David says there's no special port.

JAAP AKKERHUIS: Yes. As an implementer of this, it just uses standard ports, things like that. The extra things you need to do is that especially when it's not understood, some other authoritative serves go to loops making things difficult. And so in the end you have to fall back to non-minimization. There's a report about it by Nominet Labs a couple years ago which tells you all the troubles we had getting this done right. So if people are interested, they can find it on our website.

FRED BAKER: Thank you very much.

EBERHARD LISSE: Any other questions? Okay, Fred, thank you very much for doing this in the middle of the night.

KATHY SCHNITT: Eberhard, I think we have a question here in the back.

EBERHARD LISSE: Okay.

PETER LOWE: Hi. Are there any downsides to enabling QNAME minimization? Any pitfalls?

FRED BAKER: I don't see any. And the information that you're going to [inaudible]. But I don't believe that there's any additional problem with that.

PETER LOWE: Okay, thank you.

EBERHARD LISSE: Last question from Jaap.

JAAP AKKERHUIS: Well, to react on that, the downside is that people looking at DNS traffic and want to do statistic of it don't get the full query anymore and root server analysis, you don't see the whole query,

you only see what's needed and not really what's been queried before. So people complain about it. Statistically, it will be less feasible, but yes, that's the whole idea anyway.

EBERHARD LISSE: Okay, thank you very much. I'm going to let you go for lunch for one hour and 15 minutes if I'm not mistaken. It is 5:15 UTC, that would be 1:15. local time. Thank you very much.

[END OF TRANSCRIPTION]