KOBE – Coming Up With Best Practices to Improve Security in the DNS Ecosystem
Wednesday, March 13, 2019 – 11:00 to 12:30 JST
ICANN64 | Kobe, Japan

UNIDENTIFIED FEMALE: Good morning, everyone. This is the Coming Up With Best Practices to Improve Security in the DNS Ecosystem session. We will be starting in a couple minutes. Thank you.

PAUL HOFFMAN: So—

UNIDENTIFIED FEMALE: [inaudible]

PAUL HOFFMAN: Oh, do you want … Welcome to the meeting. Feel free to come down further if you want. So today's session is one of the first of many – actually it's not the first. We've had other ones in the past few days. But this is the first general session on talking about the recent DNS security incidents that happened, and not just talking about them but hopefully starting to come up with some best practices to help prevent them in the future.

Here's the agenda. So I'll go through this quickly. I will talk a little bit about why is ICANN [inaudible] this, although I would hope

that'd be obvious. Merike Kaeo will give an overview of the ecosystem and some of the normal best practices that we're already seeing.

Then Tim April will talk about the actual recent attacks, how they were mounted, how they were discovered, as a way of diving deeper into the kinds of things that you might except to be happening more in the future. We want to emphasize that that discussion is not supposed to be limited to, "Okay. This is the entire set of attacks, and if I'm impervious to them, I'm good." This is an indication of how good the attacks are getting these days. And Danny might speak with some of that as well.

Then, after that, Harald will not do what I said on the bullet here but will talk a little bit about the Board and the Board's concern with this.

Then we will, unless something very badly goes wrong, definitely have time for questions at the end. So questions and comments, although questions are always more useful because they're usually more useful for the rest of the people in the audience. But we should have time at the end, and we'll break by 12:30 so you can get lunch.

With that, let me just do a little bit of – oh, I'm sorry. I didn't introduce myself. I'm Paul Hoffman from ICANN org. I'm in the Office of the CTO. So why are these topics interesting? Well,

hopefully that's pretty clear, but really the most important thing that we've been seeing is that the recent set of attacks that have been noticed are coming from people who want to interrupt the DNS by inserting their own information. We're not talking about DDoS attacks here today, which is one of the things we have often talked about. We're talking about players who want to change some of the DNS information to their benefit.

How do they do it? Interceding communication or they take over a name server directly. The reasons for them doing this in the past have been fairly obvious. It's getting more subtle now. Especially when Tim talks, you will see that the way that they do it is meant to keep the attacks secret.

But given all of that, we want to find ways to prevent that from happening for the benefit of the DNS ecosystem. Merike will especially be talking about some of the current practices, but we hope that you all will be thinking about practices that you need to implement; if you're a registrant, protecting your own; if you are registrar, helping protect your customers; if you're a registry, helping you protect the registrants and the registrars. It is an ecosystem, so ICANN wants us to deal with this as an ecosystem problem.

And we want to come up with best practices. So for those of you who aren't able to spend a lot of your time on security, you can still have more of a secure system.

So with that, I'm going to turn it over to Merike. If we could have the next slide deck, that would be great.

MERIKE KAEO: Thanks, Paul. This is Merike Kaeo, and I'm speaking as a member of the community and somebody that has been doing security for over 20 years. So as a we look at the overarching ecosystem – I will just use the global Internet, not just the DNS ecosystem. When we start looking at what the most common threats and attacks are, I always classify them into five specific areas.

One – this is quite critical – is any unauthorized access. So either through vulnerabilities through insecure hosts or some kind of a password cracking or credential cracking mechanism. Also eavesdropping – being able to actually sniff on the wire, looking for any kind of information that then somebody can use for some kind of attack, malicious activity. Hijacking – Tim will talk about registration hijacking – is taking over communication. A lot of times what happens is that you're modifying data to impersonate somebody. Also spoofing – faking and impersonating somebody to fool access to control mechanisms or redirecting traffic. Then also of course DDoS attacks.

One of the problems that I always see in security overall is that people like to use different terms and make up terms. Certainly, vendors are really good at making up new terms. We really need to always get back to the basics in terms of what are we talking about with threats or attacks or mitigation techniques.

Over the last ten years, DNS abuse has been rising a lot. These are just some highlights from things that I took a couple years ago. But even more recently in the last year or two, there's been even more sophisticated attacks against the DNS ecosystem and infrastructure. I really do think that, as a community and as an ecosystem, we have to start paying attention.

So what is actually changing? The answer is really nothing. So when you're looking at security controls, you always have to think about these six fundamentals: user and device authentication, user and device authorization, data integrity – data integrity is big. Usually the mechanisms make sure that you know who you're talking to. So there's data origin authentication as well as integrity in terms of that the communication hasn't been modified in transit. You've got a confidentiality aspect. While important, personally I always think the data integrity is much more important so that you know that the data that you're receiving is coming from who you're expecting and that nothing has been changed in transit.

Auditing and logging. It's very important. People do it very poorly because often you've got some alerting mechanisms. If you're actually looking at infrastructure changes, we can say, "Hmm. I authorized that change," or, "Is there something nefarious going on?" Then there's of course DDoS mitigation.

So the security controls and what you should be thinking about is the same thing that everybody's been discussing for 20 years. But what is really changing is the scale of attacks, the automation, where people are just scripting code and putting it up on GitHub so it's easy to have malicious malware that's automated, and the sophistication. The sophistication is something that we really need to pay attention to and be more concerned about.

So as we look at the DNS ecosystem, we need to think about people and processes. We need to think about ICANN's role, the registries, registrars, reseller communities, the registrants, and everybody else that's part of the issue. We'll talk a lot about credential management because that really highlights a lot of the fundamental means that a lot of these attacks are successful at, and that talks about people and process.

But if we look at what are the technical threats against the DNS ecosystem, they can be quite vast. So I received permission to use this slide because I really like that you start categorizing all the different types of attacks that you can have against the DNS.

For the most part, I always say people try as best that they can. I believe in good intentions. No matter how good you are, you'll still screw up because we're humans. You're going to have protocol errors. As people in the IETF come up with protocols and people start operationalizing them, you say, "Oh, I didn't think of that. Okay, let's go back and fix that and have new implementations." If you have implementations, you have software bugs. You have vulnerabilities. Again, we're going to have them. We're going to have to live with that. But it's really important to assess the risks in your environment and look at what are the fundamental security hygiene practices you can put in place to mitigate a lot of the known attacks that can happen or mitigate the risks in your environment.

So as we look at the sophistication of attacks, this is over a year ago where some of you might have heard about the Ethereum attack. There's been other similar types of attacks happening, where it starts off as a route hijack. The [right route] hijack is such that somebody will set up and impersonate an authoritative DNS server, and really the intent is to create a cache poisoning attack to get the victim to actually go to a fake a website and then siphon off credentials or what have you.

But the miscreants, the people trying to do something bad, are now piecing together how do these protocols actually work. So it's not just infecting your end host systems with malware or

things like that but actually utilizing this critical infrastructure in a way so that you can redirect traffic.

Another example is some of you may have heard of DNS Changer, where there is a criminal community. What they did was change the DNS settings in a resolver. There's also another very large-scale attack that happened in Latin America quite a few years ago, where there were hundreds of millions of devices that were actually exploiting a vulnerability in home router. They were resetting the DNS servers, the resolvers, so that, when somebody tried to get to a banking site, it was actually redirecting to a page to links to malware that disabled banking protections.

So this is happening quite often. You may not read it about it always in the news, but there's quite a bit of it, where people are attacking the DNS infrastructure.

So I had this slide at first where I put in DNS everywhere, and then I decided to take it because these are common practices for basic hygiene and it doesn't matter whether or not you're doing it for DNS or anything else that's connected to the Internet. So keeping up with vulnerabilities is really difficult because, if you have a lot of different software in your environment or you have a lot of different systems, how do you track them all? But in some way, depending on your environment, large or small, you have to

figure it out because vulnerabilities is a very easy way, with all the hundreds of scannings that's going on, for you to be susceptible.

Reviewing and applying all-systems security patches. Again, this speaks to knowing what the vulnerabilities are and then applying the patches. You want to review logs for unauthorized access to systems. If I would give one piece of advice in terms of what you look at in logging, because there can be many different things that you're going to want to look at, unauthorized access attempts would be number one in my list of things to do.

I'll be talking more about credential life cycle management practices. Sometimes people just think about passwords, but there's a lot of different credentials that you have to look at and also [you have to] ensure multi-factor authentication.

So when we're looking at DNS basic hygiene – I'm not going to read all of these – these all [deem down] to: how do you have best practices so that your systems are available and don't go down? It speaks to geographic diversity. It speaks to operating diversity. It speaks to having enough bandwidth in place, enough computing power, so you don't get a DDoS.

Then, additional basic hygiene also determines: is it a need-to-know basis? Why do you have an open recursive resolver that's unmanaged? That means that can then be utilized to do so harm.

So you really need to think about engineering your systems. It's all about control over who has access and making sure that people don't have unauthorized access or it's not easy to get unauthorized access.

So, within access, there are a number of advisories that talk about best practices. Now, as a practitioner, I also see that it's very easy for people t say, "You know, the best current practices are all known." Well, [yes], to the people that have been doing it for 10 or 20 years. But if somebody is trying to find a document that will tell you what is the best practice, the cool thing is there's so many to choose from. So I really like the initiative by ICANN org to actually look at, in this day and age today now, what are the DNS best practices that can actually be – it's never easy – operationalized.

So let me spend a little bit of time talking about credential management because many, many of the successful attacks start with somebody being able to impersonate somebody because a credential has been hacked. I say "best practices for credential management" because there's just so many best practices to choose from. But there are at least two or three items that you can do that everybody knows that really everybody should try to operationalize.

So there's many types of credentials, and I started using the word "credential" probably ten years ago because, usually, people have good practices for certificate management, but not for passwords, not for security tokens. And biometrics is also starting to become a reality.

I'm not going to go through these next two charts, but I wanted to do was just emphasize the fact that it's not just one password that you have to keep a track of. In your environments, in registry and registrar operations, you have multiple different systems that you're probably logging into. I really hope that you don't use one credential for everything across the board. Good practices also do try and minimize fate sharing. This particular chart is also listed in the SSAC Advisory's 074, which the SSAC specifically came up with four years ago to raise the importance of paying attention to credentials and making sure that you have effective credential management life cycle processes in place.

And think about it in your personal environments. I actually have an encrypted spreadsheet where I have all of my accounts and all of my passwords. I have 180 entries. Think about it. I'm sure you guys do, too, especially the women who do online shopping, where you have to have an account before you can buy something. So I've created my own categories, but l literally have 180 different entries and credentials. It's crazy.

But if you're thinking about your business environments, I can assure you you probably have three or four or five if not more. People really need to pay attention to, "Well, how do I make sure that these are not compromised?" So this is just a picture that shows the interaction between the DNS ecosystem. But it's very simplistic, obviously. Basically, everybody is logging into a lot of systems, and you have to think about who has access to what, given, if you're a registrar trying to get access to a registry system by [subversion,] you do whatever you need to in your operational environments.

So credentials are compromised on a daily basis. So most common is being a victim of a phishing attack or also all the breaches. At this point in time, the reason that I have such a strong view of that you have to do multi-factor authentication is because I think almost every single company in every single industry has been breached, which means that your username and password is somewhere out there in the underground economy.

Laptops get stolen. Sharing your passwords with another person. I have had two incidents where I'm in an elevator at a conference and behind me two people are saying, "Oh, yeah. I've got to get into the system." "Oh, yeah. I don't know. Can you log in?" And they just [share] their username/password. I turn around and see what their badge is and I'm wondering what their company is.

Gosh, I can look up what their AS number is and what their [ARIN] allocation or RIR allocation is, and maybe if I was a miscreant and it was a company I was interested in, I could easily try and figure out if I can scan some devices and use the username/password. But people do this. They do this on airplanes, too. I've heard this. It's kind of crazy.

Reusing the same password on many systems. I'm just hoping that nobody who works on a registry or registrar uses an application where they say, "Yeah, I'm going to log in with my Facebook credentials," and it's also the same password that they use internally to access some internal systems. It's been done, but it's something that we need to be really careful of. It's easier for me as a user so I only have to remember one password, but really you shouldn't do so.

Anyway, there's multiple different ways how credentials get compromised. It's a very, very easy attack vector.

So this just shows in one slide things that you need to keep track of and really pay attention to. Again, I will refer to the SSAC SAC074 because we have very detailed wording in there in terms of what to think about.

Really, you want to avoid surprises. So the next slides just have a couple of points where I think everybody should look to see whether or not within their credential management process

they're being effective. So check to see whether system log passwords included a text on any authentication attempts. Make sure that you're actually using encryption for that. Some systems also sometimes have configuration files that store passwords or share secrets in clear text.

I'm not surprised by any of this. People make mistakes or they don't think about it. These are things to think about. Just look and make sure you don't have this in your environment. Make sure you know how backups are done and how a credential is stored in backups. If you're using cloud storage, just make it's not all in clear text. Some people still do have that.

These are just questions to consider in your environment. If you don't have the capability of having multi-factor authentication, why is that? Is it because you don't have time to do it? Is it because you don't have money to do it? Or is it because it'll take you some time to figure it out and, well, you'll get to it eventually? That seems to be the case. I'm a huge proponent of we in the community, if we've done this, sharing with the community.

Just know also how to get rid of credentials and accounts. So sometimes people keep older accounts around, and they can be misused.

So one of the things also – the last point is the one I want to emphasize – is what kind of measures do you employ to detect

compromised credentials? Is it just when you have a breach, or do you actually do something beforehand? This gets back to the point that I was making earlier. My single most recommendation would be, if you're looking at logs and auditing anything, make sure that you're logging for unauthorized access attempts because I think that will go a long way.

So credential management best practices just in a nutshell from my perspective. Know all of the credentials that are utilized to get into your critical systems. Limit fate sharing. Encourage the use of multi-factor authentication. Do not send or store credentials in clear text. Create effective processes for credential changes and know where you're storing them. Multi-factor is important. Tim will touch on some of them so I won't go through them.

I will point to a registry strategy from Brazil because ICANN org is starting to create a wiki where it's going to consolidate best practices. This particular presentation came from 2015 from a LACNIC TLD meeting, which the Brazil registry showed how they implemented multi-factor and some of the gotchas that they ran into. Very useful and very nice of them to actually donate them to the community.

Registry lock is also something useful to really consider because it gives an added measure of protection to make sure that you are

actually who you say you are before the registry will actually make some changes.

Then, lastly, even though it's not necessarily based on credential management, auditing and monitoring your systems also is very important.

So this is just more of a generality overall. I'm going to now toss it over to Tim, who will give specifics about the recent attacks.

PAUL HOFFMAN: Thank you, Merike. So next slide deck, please?

Okay, great. So – actually, Tim, you introduce yourself.

TIM APRIL: I'm Tim April. I'm a member of the SSAC as a [inaudible], but I'm here representing my role in the community. I was not an immediate instant responder, but I know many people that were directly involved with this response and was answering questions pretty much throughout the entire response and helping with some of the digging into the details of it.

So, this got smaller in transition. But just a quick overview of what we will and will not cover in this. We're going to try to go into the technical details that are public information about the attack. There are some things that I'm not going to talk about, partly

because this is still going on. We'll go over some of the methods used in the attack, but I will not get into any of the attribution. I usually do not. I stay away from attribution. I leave that to the professionals. And I'm not going to talk about any specific targets – not names of companies, names of countries – anything like that. Some of those are already available online. I can't speak to how accurate they are.

So at a high level, this attack was a highly skilled adversary that was able to take advantage of the transitive trust relationships that exist in the DNS. There was a very large array of targets that were involved, and the adversary used what ranged from very trivial skills all the way up through some very – they had a very deep understanding of the environment that's involved in the DNS and the registration system.

A lot of this was done through unauthorized access using stolen credentials, and many of the attacks started with changing the authoritative name servers that were used that the DNS pointed to for domains that were registered or were actually changing records that were in some hosted zones.

The ultimate goal or the perceived ultimate goal at this point was to conduct man-in-the-middle attacks on specific targets, so high-value assets within some of these zones that were used.

So to dig a little bit deeper into it, they used the stolen credentials and they were modifying either name server delegations or actual entities or records within zones that were not at the registry level but at the domain level hosted through some third-party hosting services.

While they were doing this, the common things that were changed were MS records and A records at the root or in the child zone. Then they also, in some cases, modified the DS records to modify how the DNSSEC verification was done for some zones.

Once the attacker was able to control where the name server was pointing, they could have free reign of what to do with the zone. In past attacks that were fairly common, the attackers would gain control of a zone and just not respond to it or change all of the records to point at a single host. So it was very obvious to any of the people accessing the site that the attack was going on because there was a DOS.

In this case, the attacks were not attempting to DOS anything where they would specifically redirect certain queries. So it was a very precise method where, in the normal case, they would redirect the valid queries to the legitimate name server, get a valid response, which could still be DNSSEC-signed, and return it to the user. When they wanted to do an interception, they would

respond with an IP address for the server that they were looking for.

If the zone was DNSSEC-signed and the end user was using the validating resolver, then the site would be DOSed because the validation wouldn't check out. This could result in end users noticing and going and changing something, but may it also be users thinking that the Internet is not working properly right now so they'll just ignore it and go on their way.

With full control of the name servers for the zone, this also meant that the attackers could request certificates using the domain control validation method. So there are many free services out there where you can request certificates and get a certificate for any site that you can prove that you own control of the domain. So this can give you the little green lock in your browser if you want.

One of the common protocols that was man-in-the-middled during this attack was IMAP or web-based e-mail protocols or VPN services. When the man-in-the-middling was done, the attacker was continuing to proxy the actual requests back to the origin infrastructure. So your mail would continue to work. The problem here is, if they have a valid certificate and they look in your devices and believe they're talking to a legitimate host, your device will attempt to log in, sending your username and

password. The proxy can then see that in the clear text while it's proxying the traffic through, log it, and then save it for later.

So you do this sort of thing. You can also do this for other services, too. You can do it for just standard HTTPS sites. Log the credentials and be on your merry way.

The attack that we're seeing was fairly short-lived. So it could be a couple minutes, a couple hours. It could be in the middle of the night. So if the alert fires, you'll look at it in the morning.

But once the attack happens, it can happen to just a small set of your traffic and you can get a wide swarth of credentials because many people are walking around with a phone in their pocket that's constantly polling for e-mail. If the server changes, it'll relog in and then all of the people in your network that are not using a validating resolver in some cases may end up logging into the system and losing full control of their credentials. Then the attack later can on can then log in and retrieve whatever services that you were using.

So, to quickly go over some of the attacks that were seen, we tried to make some basic slides showing the flow of how this worked. So in the first example of the attack, if you're using a DNS name service provider. So if you're not hosting your own zone, you subscribe to a service where you register the zone with your

registrar and point it at some other entity's name servers that host the zone for you.

Credential reuse, as Merike was talking about, is rampant throughout the Internet, where, if you reuse your password somewhere else, there are databased where you can go and say, "I have this e-mail address." You go look it up. You try and find the credentials that that person has used at other websites that have been compromised and try them on the DNS service provider.

If you can get in, then you can get full access to their name server. There's no registration hijacking in this case, but you can change whatever records you would like and then forward any traffic you want to the attackers' servers.

The next type of attack here is focusing on the actual registration infrastructure. So one of the slides that Merike was just showing where there was all those interconnected pieces where the domain owner, the registrant, logs into the registrar, the registrar logs into the registry, the registry pushed the zone to their name servers, and then that delegates to the authoritative name server for the zone, and then that will give you the records for the actual application servers on the other end.

The same sort of credential reuse attacks can happen here, where the attacker can somehow phish or capture the credentials for a registrant's account and then they can log in and make changes

as if they are the registrant. Those then will propagate through the backend service to the name servers and then redirect traffic to some other nameserver and application servers as needed.

This could also be if there is a super user or a privileged user within the registrar that has reuse or phish credentials or some other way of stealing them. The attack can use those credentials to make changes to their controlled zones at the registry.

So going the next level up, the attacker can phish credentials from the registrar and go directly to the registry. So this is often where the attackers could be speaking EPP or whatever other registration protocol that the registry speaks and make changes directly at the registry level.

Like with the registrar, they could also phish registry credentials if there's a super user there. They can make changes in the wider swarth of the zone, like if you were to try and phish a support user at the registry, you can have a wider impact on the TLD.

Then the same sort of thing that we were talking about earlier, where they are redirecting the traffic to the attacker's server and proxying it.

Here's an example of how that traffic interception works, where the client talks to the DNS and gets some IP address or IPv4 or IPv6 address to talk to the application server and it sends all the

traffic across there. The application server in this case probably has a TLS certificate that's issued by some certificate authority, and the traffic goes as normal.

When the attack comes in, they modify the DNS response. They point them at whatever server they want them to, and then the attacker can forward that proxy traffic back to the application server or it can deny that traffic at the proxy if it wishes to.

So this is a transparent proxy. The end users probably don't know it's happening unless they're really digging into the logs of what's actually going on on their devices. The application server can, in the proxy case, see a whole bunch of interesting data: if all of their users start coming through the same IP address, that's probably a sign that something fishy is going on.

So how was this found? It was found by multiple organizations over the course of the active attack, but the stories that I've heard about how it came to be was often through passive DNS monitoring. Using passive DNS to monitor for this sort of thing isn't always going to show this sort of attack happening because there are ways that you could mask that sort of data in passive DNS. It wasn't obvious at the time that something bad was happening. It's just it was a fishy thing that someone dug into and then uncovered this hornet's nest.

So on the topic of monitoring, these hijacks were short-lived. So you're monitoring if you're running anything for your zone, which you probably should. It has to be consistent and notice changes very quickly.

So, as I was saying, mobile devices are common for speaking to the services very often, so a short-lived attack can be result in a lot of high-value assets being compromised.

If you're a registry or a registrar or publishing zones on a regular basis, you need to be aware that the attackers can learn most of the timing constants that are used in your network. So if you're relying on only publishing every five minutes or something like that, that can be noted and probably exploited further on down the road.

The monitoring that you do run should always store any error conditions it sees because, if this sort of attack happens at 2 A.M. and everyone in your company is asleep, you want to be able to go back and look at the data to see what actually happened because, when you wake up, the problem may be gone.

Then there's a tool that Verisign has published online that lets you go and look at some of these transitive trust issues that may exist in the DNS. It's not a comprehensive tool for all of the trust issues, but it's what's happening in the actual resolution process for domains.

So here's an example. I ran just dot through it. You can see that it's kind of fuzzy but there's, in the center, just the name dot, and then you can see that it has all of the root servers. So all 13 root servers are listed and directly connected to rootservers.net. So they're all kind of in the same trust domain, but you have to notice that each of these root servers are operated by different entities.

If we go one level down – I picked out JP because it's where we are right now – you can see that the rely on the same trust circle for all of the root servers. Then, off to the other side, all of the JP naming infrastructure.

So this is actually a fairly clean instance where there's two separate silos. There's one interconnection because the room has to delegate to .jp. But, if you look at some other TLDs, you'll see a much less clean picture, like .org, where you see the root name server down at the bottom corner and then there's all of the other .org infrastructure that is all over the place. I believe from looking at it earlier, most of that was all controlled by the situation or two organizations, but there's a lot of other servers that are involved in serving that zone.

Then, because [we're] ICANN, I ran ICANN.org through. You can see that the map just gets bigger and bigger as you go out. The tricky part is you only have to compromise one of these machines

to get any traffic. So you're putting a lot of trust in the entire network.

Oops, wrong button. So, yeah, there's ICANN.org from that previous slide.

So, as we were talking about, monitoring is key. There's no silver bullet for how to detect any of these sort of things happening. Monitor your infrastructure and the zones that you host your DNS on. Monitor your parent zone. Monitor anything that's in that transitive trust graph that may change how the delegation happens. Monitoring diverse geographic places may also provide interesting evidence is there are localized attacks.

As I was talking about earlier, certificates were issued in the recent attacks. The certificate transparency logs were updated during them. So if you do monitor that, you may have been able to catch that sort of thing if you were a target.

DNSSEC validations. If you monitor for that on your own network, you may see interesting things that are going on there. If you are a registry or a registrar, monitor authentication issues within your EPP environment. I'll get to more suggestions about EPP in a second. Monitor your nameserver records. If you run any sort of network, have a domain, make sure your contact information is updated. And monitor for any changes there because sometimes those sorts of records can be changed and you can actually

request higher-value certificates, like EV or OV, using change records there.

So I mentioned DNSSEC a couple times. If you run a zone, strongly consider signing your zone. It can prevent some sorts of attacks, but it's not a perfect solution. If will help identify integrity sorts of attacks. So, in the case your zone was signed in this proxy condition and they didn't remove the DS record from the root zone, any user using a validating resolver would have not been impacted by this, which brings up the next point of: use validating resolvers for all your clients and services. So this may be you're using some public service that's doing validation, or you may configure validation for your internal resolvers.

When you are traveling, that's often a case where you may end up getting downgraded from validating resolvers to non-validating resolvers, which can open you to other types of attacks.

Moving onto EPP security, EPP sometimes isn't looked at as carefully as the other protocols on the Internet. So, if you're running an EPP service or you're a consumer of an EPP service, make sure you're using secure transport for all the communications. The passwords that traverse EPP can be used by other people to then do whatever they want with your account. So validate the TLS certificates on both sides of an EPP

transaction if you're using TLS mutual auth, which is also a good idea.

If you're the operator of an EPP environment, IP ACLs are useful but should never be used as a primary control for preventing this sort of attacking. It's a helpful secondary control but is not going to prevent everything. And of course, as we were saying earlier, make sure you log everything you can with EPP. Review them and take a look at them later. If you're a registry operator, strongly consider allowing registry locks if you don't already.

So to wrap this up, many things rely on the DNS and actors are starting to target it more often. Credential management is key, as Merike was talking about. DNSSEC can help, so please sign your zones. There's very weak adoption for DANE and CAA so far. CAA probably has more adoption by the certificate authority so far, but I strongly recommend looking into both of those systems and deploying it if you can. Monitor your infrastructure and understand the systematic dependencies within your environment.

One other piece that I couldn't find a great place to put is, in this particular event, the victim notification was very difficult for the responders. In the middle of the night, trying to contact someone's WHOIS record for their TLD didn't always work and meant that there may have been victims that were impacted for

much longer than they would have been otherwise. So update your WHOIS record. If you have security@ whatever you're in or a phone number, make sure it's monitored by people as often as possible.

So I think that was it … yeah.

PAUL HOFFMAN: Thank you. Harold, do you want to introduce yourself and tell us why you're speaking to us?

HARALD AVESTRAND: I am Harald Alvestrand. I am the IETF liaison to the ICANN Board. I'll say some words about why the ICANN Board and me personally are concerned with this matter. The ICANN Board isn't an interested party as such, but we are concerned because this is actually very concerning.

We see that there are entities that are using the DNS to attack people. Needless to say, that's not what we want the DNS to be used for. These entities do not care about the integrity of the Internet. If they damage our reputation, damage our ability to work, or damage the openness of the Internet, they don't care. But we do.

These are competent, persistent, and not going way. In some ways, they resemble elephants trampling through a spiderweb. It's the web of trust that they're trying to crush. We, the ICANN community, are stewards of that web of trust, where we want people to use the Internet for legitimate purposes and make sure that we are assured that, when they're using it for legitimate purposes, they can do so in confidence and safely.

As the talks have demonstrated, the web of trust is built on many strands that are more fragile than I want to think about. We depend on other people doing their security well, and if that doesn't scare you, I don't know what will.

We used to point to DNSSEC as one part of the solution. And it is, to one part of the problem. A number of years ago, we had a demonstrated series of attacks called cache poisoning, where it was pointed out quite correctly that this particular attack wasn't possible when DNSSEC was in place. The attacks we're seeing today are of a different nature. DNSSEC will not defend you because defending against this set of attacks requires that multiple parties have proper security hygiene focused on credential management.

When I first heard of this, I felt, "What can I do? What will I do if they're out to get me?" Also thinking about it, at least there was one thing I could do. I went to my registrar and turned two-factor

authentication for me own personal account, so at least, if someone is out to get me, in order to get to my friends or get to the things I'm responsible for, the bar is a little bit higher.

We need to take action so that the ecosystem that we are here to evolve and protect continues to be trustworthy. And that is what concerns us now: make the world a little bit safer.

PAUL HOFFMAN: Thank you, Harald. Maybe we can even make it a lot safer, but a little is a good start.

So now we have time for questions. Can I get the next slide deck, please? And can I get the clicker? Thank you.

Great. So we are ready for your questions. Any of us can answer. You didn't hear Danny McPherson speak, but he can certainly answer questions and such.

We are definitely interested in anything that any of us had said today that has sparked questions or concerns. We would love to hear from registries and registrars who have already been proactive in helping their customers. Any of that.

So, please, there's a mic line over here. For those of you who haven't been in the room, on this side there's an active mic. We would happy to answer your questions.

Great, thank you. Because I was going to start asking you questions if you didn't come to the mic. So thank you.

UNIDENTIFIED MALE: [inaudible]

PAUL HOFFMAN: But can we have this front mic turned on?

UNIDENTIFIED FEMALE: I don't know.

PAUL HOFFMAN: Maybe.

GABRIEL: I cannot… it's live now.

PAUL HOFFMAN: It is live. Thank you very much.

GABRIEL: Hi. So my name is Gabriel. Question for you, Tim. You mentioned that there were obstacles in victim notification associated with lack of readily available WHOIS information. I wonder if you could

expand upon that, especially given the WHOIS issues that are talked about in other parts of this conference.

TIM APRIL: I wasn't the one doing the notification, but the contact information that they were trying to access was not fully complete. So some of it was from the IANA database of domain registrations. They were also trying to contact some end users.

The common way that people try to reach out to some security people at organizations is the security@ whatever domain you're attacking about. And that doesn't always work.

So – yeah, go for it.

MERIKE KAEO: I want to make a more general statement. This speaks to having appropriate incident response plans. Most of the time when you have an incident response plan, you also have a list of contacts in terms of the people that are important to you. So you don't always rely on going somewhere out of the network because your network may be down.

So it was more of a statement of making sure that you can get access to the contacts and that you actually have, in your incident

response plan, a table with all the critical contacts and make sure that you update them.

PAUL HOFFMAN: Let me actually hop in with something else that is sort of a synthesize of what Tim had spoken about earlier. Some of these attacks were short-lived, as you said, so even if the contact was good and the person who saw the attack called the contact at the domain and said, "You've been attacked. Someone changed your name server records, and this happened," that person might look at their service records, if they don't have any logs, and go, "No. there's no problem here."

So beyond just having good contact information, you want to have good contacts. You want somebody who, when an attack has been reported, has the ability to do due diligence to see whether the report is good because effective attacks will be obscuring and confusing, as Tim said, because it was quite short-lived. You want somebody who has the ability, especially [to look in logs], to say, "Oh, yes. We were," as compared to, "I don't see it. Therefore, I'm not interested."

GABRIEL: Thank you.

HAROLD [BEN]: Hello. It's Harold [Ben] from Samsung Electronics. I'm actually here representing Etsy. So it's more of a comment, really. On cybersecurity, Etsy has got a whole raft of specifications and guidelines. So I'm sure the community would be welcome to download their available freely to anybody on Etsy.org. If you just type in "cyber" in the search bar, you'll get a whole list of them.

The reason why it's becoming more important for us and in fact DNS security for the mobile community is that our networks are moving from having very proprietary, very locked, tightened infrastructure to a more open Internet model, and our NFE network function virtualization architecture that we're having now relies on DNS. So it'll be very interesting to see in the future how secure our mobile networks become because the consequences of somebody getting into a mobile network can be horrific. Thank you.

PAUL HOFFMAN: Please.

JAY DALEY: Hello. My name is Jay Daley. Nice to meet you all. I was talking to a small ccTLD earlier about this concept of best practice for registrars and how they secure their interfaces. They pointed out to me that they have a number of small registrars, really quite

high quality ones, that do everything manually. So, for example, there's a specialist Christian one that has about 300 customers but does everything manually.

How do we develop best practice that scales from something that size up to some of the huge people in our industry?

MERIKE KAEO: I would love to see tools that had some basic defaults that offer better out-of-the-box security and also ways to automate. So, from this ecosystem, looking at the tools and the processes that these smaller players use, it'd be useful to understand where improvements can be made for automation and potentially also defaults that might help with adding security, if you will, out of the box.

PAUL HOFFMAN: So I will pop in. I'm almost on the other end of that questions, which is: [from] a small specialized registrar – and there are actually many of those in different parts of the world – when presented with the idea that they maybe should offer two-factor authentication to their registrants, what I have heard over time is, "Oh, that's overkill." From what you've seen today, no, it's not.

It would be interesting if the ICANN community, especially the registrar community, could come together and work on some

guidelines so that even a very small registrar would feel that this is something that, if they wanted to offer to their customers – we know that there'll certainly be registrars who don't care about this, but as you were saying, some of these small registrars are small registrars because they actually do care about their customers – they would be able to go from today's "You pick a password" – if you're lucky, you're told it's supposed to have numbers or whatever – to much better credential authentication for their registrants.

At the same time, it would be nice if small registrars – of course, we know from ICANN there's a wide range of sizes of registrars – could do that for EPP as well.

So, please, if you had a follow-on.

JAY DALEY: Yeah. So the type of small registrars I'm thinking of manually process every request. So they eyeball everything. So that's where the difficulty comes because I think that they don't need multi-factor authentication. They know the people they're dealing with, and they know their habits and other things. So they're like an AI system but now A on the front of it.

PAUL HOFFMAN: With no A there.

JAY DALEY: Yeah, that's right.

PAUL HOFFMAN: It's an RI system.

JAY DALEY: That's right. Exactly. So that's the complication that I see. That's one example, but I'm sure there are other examples out there.

PAUL HOFFMAN: Well, the next step up from somebody who processes everything manually would be somebody who used to process everything manually but it got a little bit too big. How do they go to the next step of automated processing and go all the way to the step of doing it correctly the first time, not saying, "Well, we'll just start with passwords and then eventually we'll get to the good stuff"?

JAY DALEY: Right. I have another question.

PAUL HOFFMAN: Sure, please.

JAY DALEY: Holding up the people behind me.

PAUL HOFFMAN: Right. That's okay. You can cut in front of them.

JAY DALEY: Great. Thank you. So the best practice that we've heard today is excellent. Obviously some of these people know something about it. How do you think we take this form of best practice to get to some form of standard that registrars can pick up and say, "Right. This is what I'm going to attempt to implement, and I will self-certify against it or say that I've delivered to that"?

[DANNY MCPHERSON]: One of the challenges that Merike pointed out that it's which best practices. I think one of the things that's a little frustrating is there's a lot of SSAC advice, for example, and other advice far broader in the community about good hygiene and what good looks like.

So I do think that there's probably some communication stuff. Actually, I think that's one of the steps that ICANN took with one of their first communications related to this activity and said, "Hey, here's the six or eight, ten things that are what good looks like related to this particular activity." There's certainly a broader

array of other resources out there. You heard from the gentleman at Samsung about the resource he was referencing, and there's the Center for Internet Security's top 20 control sets and so forth. There's a lot of other things that are out there that the folks should be looking at for what constitutes good cyber hygiene and what aspect or profile of that and where they sit in the ecosystem should be applicable.

Quite frankly, [in] some of the controls in a few places in the ease and broad array of attacks, they spoke about how folks should have considered some broader monitoring or some more credentials or multi-factor auth or registry locks or other things. But at the same time, it's that hindsight function.

So I do think that there's probably a role for ICANN here and more broad communication. I don't which aspect of ICANN. I don't know if it's ICANN org, if it's the ACs and so forth. But a lot of the recommendations and the [primitives] are there. Just how do you package those up so they're most easily consumed by a registrant or a registrar or a registry?

JAY DALEY: When you say "consumed," do you think that includes training for them as well?

PAUL HOFFMAN: Sure. Training, documents. I was just going to ask them, but now I'll ask them on mic. So you three SSAC people, if you were going to pick one SSAC document that people should start with, is it SAC 074?

UNIDENTIFIED MALE: [inaudible]

UNIDENTIFIED MALE: You're the expert on that one.

PAUL HOFFMAN: Yeah. Right.

MERIKE KAEO: Well, I would say yes because I helped drive that work and I'm a firm believer of it.

Also, I want to mention that ICANN org is executing on that, where they are creating a training. The first one will be at the GDD Summit in May in Bangkok. So there are already looking at how do you do the training, what exactly should be done? When I pointed out how to do multi-factor authentication, that came from the registry in Brazil. ICANN org is on its public wiki trying to

get community input on best practices in terms of what they did – "Here's what they did. Here's how we did it. Here's what didn't work. Here's some gotchas" – to help each other from a community perspective.

My perspective is that, because credential management is such a key issue for how a lot of attacks are being realized, that is good focal point to start with.

[PAUL HOFFMAN]: Good.

[DANNY MCPHERSON]: Let me just add really quickly as well that SAC 040 and 044 and 074 were kind of the key ones. Those are a decade old and—

PAUL HOFFMAN: So for those – again, you speak fast. SAC 440, SAC 44, and 074.

[DANNY MCPHERSON]: Yeah. That's the collection of those three with regards to multi-factor authentication and credential management at a registrant to registry level and registrar-registry level. There's some good basic recommendations and a lot of references in those. If you're in this community and aren't familiar with those, you probably want to be. Quite frankly, I think that at some point you're not

going to have any excuse after this most recent campaign related to these activities. But I would say that these are also a decade old and a big part of what good looks like. I think that's why ICANN [commuted] that and why SSAC put some of that work together in the past as well.

PAUL HOFFMAN: For folks in the audience who are now looking to finding SAC 040 and 044 and 074, as you're reading them, if you find questions, if you find holes that you're concerned with, certainly get in touch with SSAC. They would like to know what's going on in these.

MERIKE KAEO: Get in touch with ICANN.

PAUL HOFFMAN: Or you can get in touch with ICANN org, but quite frankly, you folks have written the documents better than we have. We have about two and you have over 100.

So thanks. Next one?

PETER KOCH: Peter Koch, DENIC. I would like to raise – well, I should [save] my words, maybe – two points. I'll start with the first and then cede the microphone to the next person.

So I think this is also a bit about focus and attention. I think I heard Harald say that, when you saw this, you found out that while DNSSEC wouldn't help, it needs something else, some other things – credential management and so on and so forth.

So does it mean ICANN, because DNSSEC has – well, a not-too-ideal deployment by now. Should people forget about it and just focus on the next thing, or could you clarify that so that we don't run in different directions every other week? That's also a bit confusing for people, I would think.

HARALD ALVESTRAND: So, since I spoke a little bit unclearly on this point, what's been pointed out is that the web of trust has many strands in it. In the case of the cache poisoning attacks, one of the strands seemed to be extremely thin. DNSSEC is replacing that particular strand with a steel cable, so you can stop worrying about. I like very much the idea that there are problems you can stop worrying about. I think, having worked with security technology in many ways for many years, making sure that there are pieces so you can stop worrying about this is a critical part of getting to a system that you can worry less about as a total.

Yes, DNSSEC deployment is somewhat below what we expected ten years ago, twenty years ago, but I think it's about what we expected five years ago. So they might be getting there.

[DANNY MCPHERSON]: So, if I could add to this, I think that this is where one of the key jobs of anyone in an operational security or an operational role is to understand what you care about and where it lives. If you looked at the dependency graphs stuff, for example, that Tim put up in his slides that he and I put together, then one of the interesting things about that is you might have 13 root servers that are Anycasted in over 1,000 locations. But if you have DNSSEC, then, if you sign the records that are distributed out there in the zones, you can detect failures or tampering of the content of that. If you compromise any one of those other servers and somebody is doing validation, only one server at the end of day ultimately matters: the one that's doing the signing and so forth.

So it adds a lot of resiliency. You minimize your attack surface considerably with DNSSEC from that perspective. So I think that's an important aspect from DNSSEC.

I think the other thing that some of the parties have looked at is, while the attackers did have access in some of these scenarios to the configuration panel at a registrar, they didn't necessarily much with a DS record or with other contents. So, if you're monitoring either changes to those records or if you're doing DNSSEC validation, then it certainly helps you to detect. It's

certainly not a panacea, but it definitely allows you to have more eyes on this. I'm all about layered security and I think that it definitely providers another layer of capability, even at the authoritative level in these examples of attacks as opposed to something like cache poisoning.

DREW BAGLEY: Hi. Drew Bagley from the leadership of the Competition, Consumer Choice, and Consumer Trust Review team. I just wanted to applaud your presentation today and note that one of the recommendations that we directed to the ICANN Board, Recommendation 14, actually calls for the incentivization of best practices against DNS abuse and stated broadly the purpose: so that it could encompass things such as the issues discussed today. So what we're hoping for is that, in the future, we would see some sort of incentives, whether they're financial or not, baked into the agreements themselves to get registrars and registries to adopt best practices.

So I just wanted to note that that's part of something that's already an issue to the Board. Thanks.

PAUL HOFFMAN: Great. Thank you.

[YOR MCCORAN]: Hi. [Yor McCoran] from ISOC-[IL], the ccTLD manager for .[il]. Two questions regarding regulation. Is ICANN considering to use its kind of regulation power with regard to enforcing all kinds of practices. One aspect.

And what is the relation to national cybersecurity regulation? Is it supposed to be kind of interaction between them?

HARALD ALVESTRAND: Channeling Goran, ICANN is not the regulator. It doesn't claim to be. Doesn't want to be. Doesn't want to behave like one. But we certainly want to encourage good practices, but ICANN is a private sector organization that has contracts.

With regards to the relation to the cybersecurity law, the law is the law, and that's a completely different sphere of control or influence. We have learned, for very good reasons, that getting things done by passing a law often takes time and has unintended consequences. So sometimes it's better to just get the job done.

But sometimes – now I'm speaking without thinking too much first, so the virtual board sitting behind my back will have to forgive me if I mess up again – using the law and the force of the law is the only way in which you can get things done. Among other things, the law is the only effective way to actually punish criminals.

So we should think of these things as pieces of the puzzle. When we can get them to act in harmony, we should be very happy. We can't expect that, but we can hope for that.

[YOR MCCORAN]: I know that ICANN is contractually-based, but you have basically a status of accredited registrar, right? So basically you can agree with some aspects of cybersecurity regulation in the contracts. So this is for the ICANN part. And this is not regulation comes from a state law. It comes from the contractual basis.

I think that ICANN's role is to – as to the national regulation, my view is that most of the regulators do not understand DNS and do not understand the threats in the DNS. So there should be a discussion between the national regulators and the DNS experts in order to promote those aspects because, in the end, these are entities that work under a specific legal framework.

PAUL HOFFMAN: So let me jump on these and then, John, you're supposed to get to the mic now because I'm about to call on you, John.

So for the first part, I'm going to be picky about your language. You said, "Will ICANN do something?" My picky answer is we might, we might not. It sounded like you followed it up with "We should," and that's great. ICANN is very, very community-driven.

If the community wants us to change contracts, to add contracts, or do whatever, we will follow what the community wants.

We as ICANN org staff have no problem with hearing multiple loud voices from the community on various topics, and that's why we have processes to come to consensus. If there is consensus within any of the communities that we should be changing the way we do things for this feature, we will. We are very community-driven.

For your second question – John, you're supposed to be going to the mic now – we are doing outreach already to some of the governments and law enforcement, and the person who's getting up very slowly is the [stuckey] in ICANN org for that. I would like him to answer instead of me. Thank you, John.

JOHN: So, actually, I'm not the only [stuckey] for that, right? So interacting with governments on this, ICANN actually has a government engagement department as well that directly interacts with governments. We also have what we call the Global Stakeholder Engagement that deals with also the wider community. Goran – not Goran but the other Goran; they have very similar names – actually alluded to a document that we had put out a few days ago or last week about how ICANN should continue to interact with governments.

Now, at the operational level, when there are incidents that affect the ecosystem that we're involved in, we do indeed communicate with governments. We also work with the GAC. We've done training for the GAC members.

So there is actually a lot of work going on there, but it's not our mandate to go out and educate every single government in the world on every single cybersecurity issue. That's a kind of boiling-the-ocean issue. But when it comes down to things like this that affect the ecosystem, we are fairly active.

But I would recommend that you read the document that they put the other week on proposals of how ICANN should actually interact with governments, specially about all the regulations they're doing and how it affects the ecosystem. We've seen some recent regulation – GDPR, for example – that actually affects the ecosystem, and we're seeing a lot more. So we put a proposal out. Go read it and please comment on that.

PAUL HOFFMAN: Great. Thank you. See –oh, please. Yeah.

GEORGE MINARDOS: Hi. George Minardos from the .build registry. First off, I've found the increasing-security sessions at ICANN really helpful as a registry. I commend you on all that. I'd like to learn more.

So one key thing I heard here is the concept of standardizing registry lock. How can registries get involved? Or what sessions should I be attending, or special groups, to learn more about that and participate in it.

PAUL HOFFMAN: I don't have an answer, but I believe the person who got up behind you in line does – no? Okay.

Actually, Danny, can you do that? Great.

DANNY MCPHERSON: So, yeah, I've heard that in a number of places. Actually, two places in particular. SSAC is going to look at the topic. That doesn't mean we're going to take an action on it. But I know that we decided that we would look at that and see what might constitute doing something more with that.

There is also a request for that on – I believe it was either DNS Ops or the Registry Operations Working Group. You can find that in the archive and comment.

Paul, do you know? Do you want to comment on that?

PAUL: Yeah. The IETF is in Prague in two weeks, and Wednesday, during one of the open sessions, there is a two-hour block to specifically talk about defining what registry lock is.

[GEORGE MINARDOS]: Okay. Good.

PAUL: And it will probably wind up in Reg X, but we're not sure.

[GEORGE MINARDOS]: Okay. Appreciate that.

PAUL HOFFMAN: So, to follow up on that last comment, unfortunately, that will probably not be either recorded or broadcast, just because of the way that the Wednesday afternoon sessions go. But hopefully there'll be written output from that that will come back into the ICANN community.

GEORGE MINARDOS: Okay. I'll look for it. Second question I think might be for Danny or whoever's in charge of the Verisign Labs. Who can I have explain the trans trust input? Because I don't get the graph and I just want to learn more on that.

DANNY MCPHERSON: So there actually was a paper there. I was just looking. I was going to mention that. So we wrote a paper that was published a while back that talks about that and Web PKI and some other things.

So, yeah, I'll get your info offline. If anybody else would like to see that paper, you can probably just Google my name, Verisign, and transitive trust, and you can find that power. But there are a bunch of others that have been written on the transitive trust and web of trust challenges with DNS and [one up], but this one speaks particularly to transitive trust.

So I'm happy to talk to you about that.

GEORGE MINARDOS: Great. Thank you.

PAUL HOFFMAN: Thanks.

PETER KOCH: Peter form DENIC. Paul, when you asked if I had an answer, no, probably not, but a caveat maybe in the direction of registry lock.

First of all, whatever the IETF is going to standardize or not in the end, this will not cover the business processes behind what is going on there, and rightfully so.

Just for the record, I also don't see a role for ICANN in that regard.

Also, as food for thought, maybe diversity in operation with registry lock might actually turn out to be a security win rather than a loss.

That brings me to something else that actually Jay brought up but the panel didn't really bite. This doesn't seem to me like a technical problem. The attacks that have been executed aren't really new. They are sophisticated maybe in the way they have been executed in terms of time windows and so on and so forth. But the path has been available or known in theory for decades.

However, it is an economic problem. We have as an industry managed to get the whole registration, publication, and so on and so forth processed for the DNS, highly automated and efficient. With that, we have managed to get the prices low – well, some of us not so. Yeah, appreciate you smile, Danny.

But with registry lock, this is an interesting thing. I wonder how it can scale. I've heard voices already that say, "Oh, yeah. Let's automate registry lock." I think this platinum card approach that Jay mentioned, the premium registrar that watches at every

request, allegedly having no two-factor, do have a second factor because they know their customers. That is the second factor, just not automated.

So that approach is nice and it's available to some. It's probably not going to scale, and we should make this promise and should make everyone believe that, yeah, they are standardizing this registry lock and it will be easy and then we'll all be saved. There's more to this. This is just one thing but it's probably not the silver bullet. Thank you.

PAUL HOFFMAN: Thank you. Please.

JAY DALEY: I'll just follow up on what Peter has said. I just want to caution anyone about thinking that registry lock can either be standardized or can be adopted widespread. There's a significant concern amongst TLDs that, when you're selling a relatively low-priced product, you allow for ubiquitous take-up of that. So you allow all sorts of marginalized or otherwise dispossessed communities to be able to access that technology because it's a ten-buck-a-domain-name type of thing.

Registry lock, because of the very nature of it, can be an order of magnitude more expensive. By doing that, you end up only

making it available to a relatively small portion of your customer base: those who have the greater resources. So by definition, you're now providing security to a group who are already in a better position to be able to secure themselves.

So that's a cultural and principle-based issue that a number of TLDs around the world struggled with around registry lock, which is one of the reasons why it's not been adopted that widely.

Okay. Thanks.

DANNY MCPHERSON: I wanted to agree with both you and Peter on that aspect. I think there's a lot of business processes that vary per registry. There are probably some aspects of a registry lock function that could be commonly known and then some that are going to be business processes. But basically I agree with everything you and Peter said, but I would also acknowledge that I think the registry lock would have mitigated this attack vector completely had it been in place. And I think that, if you have critical domain names, you should strongly consider what registry lock would mean to your operating environment.

PAUL HOFFMAN: So I see the line's done. We're almost done. I want to wrap up, since we went off a little bit on registry lock – but it actually comes

back to something that Merike started us with, which is credentials and authentication, not reusing passwords – some of these very basics.

One thing that people in what I will call the password industry know – that is, websites that let you in; they want to know who you are but you have a password and they don't really care about it – have discovered over the last twenty years – is that, with "I've lost my password link," whatever happens after you click that can either be done well – and that's expensive – or it can be done poorly and an attacker can use that as a way of taking over an account.

So we have two things that we've covered today. One is what people can and should be doing now for themselves for their customers if they're registries and registrars and such like that. And we haven't come up with a solid list of best practices, but we've certainly brought up a lot of things. But the other important factor is this does come down to money. This does come down to attention span. Is a registrant required to be thinking about this all the time? Or, as Harald said, it would be nice if some of these problems were completely fixed so they can think about other things or worse attacks and such.

This is all going to have to be balanced. It's going to have to come from many places. We would love if the ICANN community got

more active on this, both in saying what and also saying what they want to know in the future.

So hopefully this has sparked some interest. Thanks very much, and have a nice lunch.

[END OF TRANSCRIPTION]