ICANN68 | Virtual Policy Forum – DNSSEC and Security WS
Monday, June 22, 2020 – 10:00 to 12:30 MYT

KATHY SCHNITT: Hello and welcome to the DNSSEC and security virtual workshop for ICANN 68. My name is Kathy and I will be one of the remote participation managers for the session. Please note that this session is being recorded and follows the ICANN expected standards of behavior as I have noted in the chat.

During the session, questions or comments submitted in chat will only be read aloud if put in the proper form as I've noted in the chat. I will read questions and comments aloud during the time set by the chair or moderators of this session. If you would like to ask your question or make your comment verbally, please raise a hand. When called upon, kindly unmute your microphone and take the floor. With that, I will hand the floor over to Dan York from Internet Society. Dan, please go ahead.

DAN YORK: Thank you. Welcome. Good morning, good afternoon, good evening wherever you are joining us for this virtual workshop. This is a workshop that's been brought to you as part of the ICANN meetings for a good number of years now as a product of the—as this group of the current program committee that has been working to bring you this next couple of hours here as we will talk about DNSSEC and the security of the DNS and the wider Internet. You'll see some of the presentations we talk about in both kinds of topics here.

My name is Dan York, I'm part of the Internet Society and I'll be speaking a little bit more about here other members of the program committee are on this call, and a few will be involved in some of the presentations that are going on here.

This program, this workshop and the activity is organized by the ICANN Security and Stability Advisory Committee, SSAC, with some additional help from the Internet Society, but this is the group that has brought this workshop for you.

Today, we're going to cover a number of different topics. We're going to begin with—I will have a brief comment after this with some of the current status of the DNSSEC and DANE deployment around the world, RPKI as well, we'll talk a little bit about that, and then we'll get into a regional panel which will have a couple of folks here talk about DNSSEC deployment in their particular regions. Kim Davies from ICANN will also give us an update on the KSK rollover. I will be back to briefly talk about a new project in the Internet Society that has a DNSSEC component. Wes Hardaker will give an update on the local root project and the work he's been doing with that. We will then have a panel discussion led by Steve Crocker which will be around DNSSEC provisioning, and we'll wrap that up with some final Q&A over this next period of time.

So that's our program for today. We look forward to working with all of you and answering your questions. As Kathy noted, we're open for your questions and please go ahead and leave them in the chat and we'll talk about that. I would also encourage you, as you think about this, we do these workshops at each of the ICANN meetings. So while we're doing it for this session here, we will also be planning on for the next ICANN meeting as well. so if you think about something related to DNSSEC or DNS security, RPKI, any of these different topics, do consider it. We'll put out a call for participation soon after this meeting to start planning for the next one.

With that, I want to talk a little bit about deployment numbers and what we're seeing around the world. So, first of all, one of the first things we like to show is that we're seeing a continued growth in the amount of DNSSEC validation happening. We're now up steadily over 25% coming from the measurements out of APNIC, Geoff Huston's measurements team has brought this up, it's been growing a good bit, and this is for validation by Internet service providers, ISPs validating the signatures in DNSSEC.

We're seeing that growing quite well. Here are some measurements that we can see with parts of Oceania and Asia having very high levels of validation for their particular regions.

We've also watched a number of the growth of DS records which is what gets put into the top level of TLDs, and this is something coming from the DNSSEC tools project. Wes Hardaker is the one who's involved with that and he'll be talking to us later on a different topic, but just for your knowledge.

And we're seeing a continued growth now where we're up at an increasingly growing number over this time, and it's been great to see this. We're also seeing some very large success happening within e-mail, within using DANE records to be able to set up secure encrypted connections between e-mail servers. This is one of the success stories we've certainly seen in the space with regards to this. So this chart, you can find this at stats.dnssec-tools.org, shows you where this is growing in this kind of space.

One of the other topics we talk about here is the RPKI, and this deals with the validation of records of routing in the routing system and BGP and this. We're seeing a continued growth. If you look at this, the number of routing origination and the number of prefixes are here, the number is continuing to grow year over—or this is from the last cycle, from the previous ICANN meeting is on the left, today it's on there, use the growth has gone from 18.6 now up to 21%, and we're continuing to see it kind of grow. This is wonderful, to see that we're growing in and around this kind of space.

This is the number of ROAs, which are route origin authorizations, which is part of RPKI. And again, we're seeing a wonderful curve of growth of seeing that continue to go up year over year as a good sign for validation of this in the routing security space.

Finally, let me just make a point about the ccTLDs that we have out there. We're continuing to see growth all around the world, heavily throughout Europe, the Americas, Asia Pacific, still growing in parts of Africa and in Latin America.

Overall, we're seeing tremendous growth in this, in the ccTLDs. We have a number of resources, some of whom are the folks involved in this call. I mentioned the DNSSEC-tools site which is a number of tools that people may have. We have a project at the Internet Society called the

Open Standards Everywhere project or OSE, which will be a topic I'll talk about in a little bit. But that is a place. We have some historical information at DNSSEC-deployment.org, and also, we have the great stats coming out of APNIC which help show this. For RPKI, I have a number of links here as well, and we'll be seeing some of that as we go through that.

With that, I am reaching the end of what I wanted to talk about here. As a reminder, this is our agenda for today, and so our next panel that we're going to have will be our regional panel. Kathy, are we ready for that?

KATHY SCHNITT: Sure are.

DAN YORK: Okay. Sounds good. With that, I will say thank you very much for joining us. Again, please do participate, ask your questions, do the things on here. Participate with us and we look forward to learning more about what you have to say throughout the session. Thank you all, and I think I'll turn it over to Jacques Latour.

JACQUES LATOUR: Good morning, good afternoon, good evening. Next is our regional panel, and we have three presentations. The first one is from Ms. Mastura Mukhtar from MYNIC.

MASTURA MUKHTAR: Hi. Good morning. Thank you to moderator, Jacques Latour. My name is Mastura from MYNIC Berhad, registry of ccTLD for Malaysia, for .my. It is an absolute pleasure to be part of today's workshop. Thanks again to ICANN for the invitation.

Here is the agenda for today, who is MYNIC, our history of DNSSEC development, journey to adopt secure e-government services via DNSSEC. Next slide, please.

This is an overview about MYNIC Berhad who is actually the agency under the Ministry of Communication and Multimedia Malaysia, and we are registry for the country code top-level domain name .my. Due to that, we are part of the Malaysia critical national information infrastructure or CNII. As we are the registry, we are also the key enablers of the digital economy ecosystem. So we focus to develop and promote the usage of .my domain among Malaysians.

We also strive to empower businesses and industry to become of the digital economy through the domain name industry, meaning that we encourage them to use the domain name and with our country code to be part of the businesses and also to know that they are from Malaysia. Next slide, please.

This is our core services. Basically, we are managing administer for eight domain names, including we providing services such as WHOIS, DNS resolution, and also, we provide value added services, .my domain dispute resolution and also sensitive domain name dispute resolution.

We are also the registrar for .my, and we are delivering for the registrars activities as well to .my customers where we provide the customer online registration for .my domain name. Also, customer care, support and resellers whereby our appointed resellers can deal directly with MYNIC if they have any clarification or question pertaining to .my domain. Next slide, please.

We are coming to the history of development of DNSSEC for .my. Basically, we have started research on DNSSEC deployment for .my, started in 2009 where the research conducted inhouse where we have formed our internal team to conduct DNSSEC research and also to develop the signer to sign the .my zone. Also, we continue with seminar and awareness on DNSSEC to the essential sector, for example, bank and government sector facing with public users. We also participated in 2009 myDNSSEC testbed and presented during the ICANN workshop. And in 2009, we also provide DNSSEC public trial to the public users.

In 2010, we continued to conduct the seminar and awareness program in order to enable, make what is DNSSEC to the essential sector. Finally, in 2011, we signed .my and .my DS record we deployed in root. Subsequently, we established the complete DNS chain of trust between root and .my zone. In Q2 2011, we full operation on DNSSEC to receive the DS record from the child domain which is for others from the registrant that they want to enable their own domain name for DNSSEC. In 2012, we deploy DNSSEC for our IDN country code top-level domain. Next slide, please.

We had a journey last year to secure e-government services through the DNSSEC whereby last year, we took an initiative to implement the DNSSEC for .gov.my domain. So the objective here is to create secure e-government services to support the national digital economy and increase public trust in the e-services that are provided by the government sector.

Since last year, because in the year 2018 we only have like three domain with DNSSEC enabled, last year, we took an initiative to work with the government agency to deploy and encourage the government to enable DNSSEC for their domain name. Next slide, please.

Okay, this is the five focusing pillars that we are focused on adopting DNSSEC for government domain name. The first one is policy and implementation, competency and capability, domain registrant, reseller or partners. Last but not least, monitoring and validation. Okay, we move on to pillar one whereby the policy and implementation. Next slide, please.

We established the collaboration last year, we had a series of meetings and engagement with the government sector to establish their collaboration in order to drive DNSSEC for government domain name. So we collaborate with the national cybersecurity agency or NACSA and Malaysia Administrative Modernization and Management Planning Unit, MAMPU, in order for them to secure the .gov.my domain name So for NACSA, just brief, policymaker for national cybersecurity, and MAMPU, policymaker for all government IT services and we are the enabler for .my domain. We are the registry for .my. Next slide, please.

As for policy and implementation, these are the challenges that with efface during the journey to deploy DNSSEC for the .gov.my domain. Among the challenges that we face, clarity on responsibilities between the policymakers, between the NACSA and MAMPU, infrastructure and technology readiness for the government to support the DNSSEC, lack of understanding on how DNSSEC works among the DNS administrator of the government sector.

How we overcame the challenges that we face? For the pillar number one, we had engage series of meetings finally. We managed to obtain the approval from the stakeholders, which is policymakers, NACSA and MAMPU, for the approval to support for the security enforcement to implement or to deploy the DNSSEC for the gov.my domain. So we are upgrading or conducting a technical refresh on the government infrastructure and technology in order for them to support DNSSEC.

So apart from that, we also conducted training. We conducted face-to-face physical, hands-on training. We provide them testing domain in real environment for them to deploy DNSSEC to gain their confidence. So they can hands on, they can get the [actual] on how we can assist them to deploy the DNSSEC for government domain. Next slide, please.

For pillar number two, competency and capability. The challenges that we had faced in order to drive this DNSSEC deployment, yes, the administrative overhead concern on the DNSSEC key management, they're not sure how the manage DNSSEC key, how to generate the key, so understand of the DNSSEC configuration, why they need to do—for instance, server for resigning for instance, and also, the key rollover for

ZSK for instance. And also, unclear SOP to manage DNSSEC. Unclear SOP meaning they're not sure where they can start in order for them to deploy or implement DNSSEC.

So, how we overcame the challenges? We provide them face-to-face as well the technical hands-on workshop to upskill the government sector DNS administrator to reduce the administrative and configuration issues and risks. We demonstrate them in real during the technical workshop. So we provide our processes and best practices in order to help the DNS administrator to manage the DNSSEC. This is to give them assurance so by implementing the DNSSEC without service disruption on the domain name.

Next, the domain registrant, part of the DNS [administrator,] we have to convene as well, we have to brief as well the domain registrant why they are encouraged to implement DNSSEC. So the challenges that we face in order to implement, the fear of domain service interruption with DNSSEC implementation, so they thought with the DNSSEC implementation, if there's any issues. No subject matter expert to consult related to DNSSEC issues or queries. Low participation from the government agencies due to ambiguity of the direction.

With that, we overcome the challenges that we face by having the enforcement from both the main stakeholders, NACSA and MAMPU to encourage more participant from the government agencies. Last year, we managed to conduct a training with the huge participation from the government agencies, we conducted awareness and technical workshop again to them. so we had a series of technical workshop to different group, different users, so the technical workshop that we conducted is from end to end process and through our best practices for deploying DNSSEC to prove it works.

So we taught them where to start, for instance, as the domain registrar, they have to inform the DNS hosting provider, or if they manage their own on the DNS server, they can do it by themselves. The next stage also on how they can submit or upload or update the DS record to MYNIC as the registry.

In order to reduce the human error while submitting the DS record, MYNIC implemented auto fetch the DS record through our selfcare management system. So with this facility, the registrant or the technical contact of the domain name are just ensure the authoritative DNS server come from the zone completely signed, so they just simply access to our selfcare management system and fetch the DS record. So our system will create it directly to the authoritative server and fetch directly the DS record for the particular domain name. This can reduce errors so no manual input to be done by the technical contact or DNS administrator. So for instance, if the DNS server, if the authoritative failure to configure with the DNSSEC, our system will throw a message to inform them that something needs to be fixed in order to deploy the DNSSEC for the particular domain.

So as of last year, we have 1040 total of government domain name with initiative of—we had series of meeting, series of training, virtual training, physical training. So we have fully enabled 500 domain name

.gov.my to date. We managed to get them configured to be onboard to deploy the gov.my domain without any disruption or errors.

Then we move on to the next pillar, number four, we have our resellers and partners. This is another challenge that we face because multi providers have different technical skill or different skill for DNSSEC deployment and administration. As you can see, all 60 resellers that we have, only six resellers support DNSSEC. So as for government domain name, it's really basically they have two camp, camp A and camp B for instance. For camp A, they have the domain name hosted by the government-appointed provider, and the remaining balance of the domain name by various providers. So this is another challenge that we face, how to influence them, how to train them in order to deploy DNSSEC for .gov.my domain. So some of the challenges impose cost if the domain registrant wish to deploy, implement DNSSEC for .gov.my domain. So, how we overcome for challenges number four? Next slide, please.

Okay, again, we provide technical workshop and training to upskill the competency, their knowledge in terms of managing the DNSSEC in terms of to get them clear the process and SOP in order to deploy DNSSEC.

As for the government sector, for the team that manage the government domain name which appointed by the government sector, we engage with them directly in order to deploy DNSSEC. Yes, for the domain name appointed and managed by the DNS administrator, by the government, most of the domain name under their custody successfully DNSSEC enabled without any disruption. So the remainder of the agency, we continue to contact them to give them awareness, to provide them a virtual technical training in order to get their participation to implement DNSSEC for the .gov.my domain under their administration.

So the resellers and partners, by having the series of training and awareness, the resellers and partners confident and comfortable on DNSSEC working mechanism, they know how to start, what to do next in order once the domain name is signed by the DNS authoritative server that's maintained by them. They know how to set up the resigning configuration in order [RRSIG] can be generated prior to the signature validity expiry.

So the last slide, before I proceed with this slide, I wish to thank to DNSViz team. if one of them is part of this training. Also wish to thank very much to Verisign for their work to develop to allow the public Internet users to use and leverage the tools that they develop. We have learned a lot from these tools. So we share these tools while we are conducting the training with our domain registrant and also our stakeholder and also our resellers. So this tool helps us a lot in order to gain their trust, their confidence in order to deploy DNSSEC. So by having the tools, by knowing how to use the tool, how to do troubleshooting by using the tools, they are very happy, they are very clear on how to troubleshoot the DNSSEC [matter.]

So these are the challenges that we face for monitoring and validation pillar. Unclear on how to validate DNSSEC being signed successfully, because they never know how to validate if the authoritative already signed, how to validate the signed zone before they can submit the DS record to registry. Lack of technical knowhow to use DNSSEC tools.

So we overcame the challenges, as highlighted. We provided training, documentation to the DNS administrator on how to use DNSSEC, DNSSEC analyzer, especially, where they can see if the broken change of trust, where they can see if the RRSIG already signed, the information of DNSSEC [inaudible] etc. so they are happy to use and now they are familiar to use those tools in order to troubleshoot and investigate if there are any issues related to the DNSSEC for their particular domain.

In addition to that, we also performed rechecks on the zone file and full chain of trust is tested before the signed zone propagated, before we send up the signed zone. So we do perform the prechecks. These are the six prechecks that we perform prior the zone out to be propagated.

So if incompliant or noncompliant is detected and require human intervention, the process will stop and we will continue with the inspection to rectify the issue or the problems.

So that's all, my presentation.

JACQUES LATOUR: Thank you, Mastura. We will hold the question at the end. Meanwhile, in the chat, you can post your question. Next up—

KATHY SCHNITT: Jacques, I'm sorry, I'm going to make an announcement. Due to some of the Zoom bombing issues, chat has been disabled for participants at this time. So at the end, we will allow for some verbal time for questions.

JACQUES LATOUR: Okay. Sounds good. Next up is Mr. Molay Ghosh from Reliance Jio with the DNSSEC deployment for network operators.

MOLAY GHOSH: Hi. This is Molay Ghosh here from Reliance Jio. I represent a service provider in India. Next slide, please. We'll be going to the agenda how Reliance Jio taught through the years for the DNSSEC, [how they came up with] DNSSEC in Jio, [inaudible] DNS hijacking. We'll talk about Reliance Jio DNS deployment, DNSSEC queries validation stats, DNSSEC deployment, impact on DNS KPI, issues seen in implementing DNSSEC in JIO network and APNIC reference data. Next slide, please.

So Reliance Jio was a service provider which came out with the mobile and broadband services in the years. We started our journey in 2010 when we acquired a license for 4G broadband license across India. Then we stated building our network and with that, we also partnered with Secure64 who is a DNS [inaudible] provider. In 2016, we started acquiring subscribers in a matter of two months, we acquired nearly about 100 million subscribers in our network.

So this is the journey so far, and in 2020, we are seeing a lot of automation deployment and also API based DNS provisioning that we're going to be implementing in this year. And in the future, we'll see full DNS signing capabilities for our own domains with the DNS signer. Next slide, please.

So, what is DNSSEC? Basically, a client which sends a query to the DNS caching server goes through the DNS authoritative server and if the response is a signed response, then the DNS caching server validates that response, verifies their digital signature, and if the digital signature does not match, it does not give the response to the client. So basically, this adds another layer of security to the client and the client does not get phishing requests. Next slide, please.

This is the DNS hierarchy in Jio network. All of our users, they query to our recursive caching server which are deployed in all the major locations. we ego and query the DNS key root servers which are like .com servers. If any of our internal domains are queried, then it goes to Jio.com servers which are our own authoritative servers, and if it queries our own domain like jio.com, then it gets a response on our DNS servers. Next slide, please.

So basically, what are the costs of DNS hijacking? DNS hijacking does a lot of damage to brand image. It also creates a lot of downtime. It creates a liability for the hijacked organization, we are liable for damage to customers. It can also do funds theft from e-mail and that can lead to company funds and also compromise that can leak confidential info. Next slide, please.

So basically, in Reliance Jio, we have currently 370 million mobile subscribers, and our DNS is deployed in close to 30 cities. We have six public authoritative servers, 158 cache servers, and two signer servers. The peak query per cache sever that we see in our network is around .35 million QPS and the total QPS across the network that we see as of today is 15.9 million. Next slide, please.

So these are the stats for our DNSSEC query validation that we see in our network. We see 94.56% of the queries validated, 4.67% of the queries are partially validated, and we see failure of around 1.8%. So DNSSEC configuration in our network was just as simple as adding two lines of configuration files to both SourceT and x86 servers. Next slide, please.

Jio has deployed close to 158 cache DNS servers configured to accept DNSSEC, and the next step is to implement the signer for all the authoritative zones. The signer will enable centralized, scalable signing capabilities for our caching servers, further protecting our subscribers from threats like DNS hijacking for Jio owned domains. Jio has enabled DNSSEC in all the cache resolvers and these are the key information. Next slide, please.

So, what are the impact that we see on DNSSEC after implementing DNS KPI? We see no adverse effect in query per second, we see no adverse effect in cache miss latency, we see no issues in CPU utilization of the DNS servers while memory utilization has increased by 2% and we see no adverse effect in cache hit ratio. Next slide, please.

So, what are the issues that we see as [happening post] implementation of DNSSEC in our caching servers? Post implementing of DNSSEC in our caching servers, we saw that some of the websites were not resolving properly. Why it was not working? Because the resolver was doing a DNS key query but not getting a response or it was getting an invalid response and the resolver was not a secure domain name and it was not resolving for the subscriber. So this was a malicious website which was not having a valid response and this is how subscribers got no response, and in turn, it was for our subscribers to have no issues for their queries.

So then we took up with the authoritative server. There was some course correction which was done at their end and now these websites are working fine. Next slide, please.

So this is the APNIC reference data for our AS 55836. This is a chart which has been available in the APNIC website for DNSSEC which shows that almost 94.56% of the queries are validated. And we started somewhere in late 2018 implementation and completed by 2019. Next slide, please.

Thank you. I will be available for any questions.

JACQUES LATOUR: Thank you. That was a very interesting presentation. Next is Yuya and Yoshikazu from GMO Internet.

YUYA NAGAI: Hello everyone. [inaudible] an example case of deploy DNSSEC by GMO Internet. My presentation will be in Japanese and translated to English by Kojima. Thank you for your attention. [inaudible]

YOSHIKAZU KOJIMA: I am Yuya Nagai in charge of the DNS service in GMO Internet. I founded djbdns in 2003, after that, I have been studying the DNSSEC since 2009. By the way, my favorite program language is Perl.

My name is Yoshikazu Kojima. I am in charge of the technical on our shared hosting service in GMO Internet. My specialty is unfortunately not DNSSEC but web and mail. Next slide, please.

What is GMO Internet? GMO Internet has the famous domain name registrar service called Onamae.com. Onamae means name. Main services, domain registration, domain registrar, DNS provider, and sell webhosting services, etc. Next slide, please.

YUYA NAGAI: [inaudible]

YOSHIKAZU KOJIMA: We started domain registrar in 1999 and in 2014, developing advanced paying services. The DNSSEC service is one paid option services. Our customer can use DNSSEC very easy, they just enable DNSSEC service from control panel, that's all. Next slide, please.

YUYA NAGAI: [inaudible].

YOSHIKAZU KOJIMA: Let me explain overview of infrastructure components. For frontend nameserver which is BIND 9 and NSD with [inaudible] [nameserver diversity.] This is because [I was stopping the service] vulnerability for single software. So [inaudible] for us to update software for security patch.

Next, for backend software, this is OpenDNSSEC. There are some solutions to support the DNSSEC, [inaudible] support inline signing [inaudible]. What is important for us is minimize changing of our existing system. On that point, ODS was the best.

I want to go into the details about ODS. With ODS and some provisioning script, we could automate from customer order to enabling DNSSEC. Next slide, please.

YUYA NAGAI: [inaudible].

YOSHIKAZU KOJIMA: Point for development, about DNS server software, if you choose from major product, you won't have program. Configuring OpenDNSSEC is very simple. If you set up with hardware HSM, probably, [inaudible] but with software HSM, it was easy. To operate this system, it should take process monitoring and back up secret key. Not so many things. To setup DNSSEC system, we didn't face serious problem, fortunately. Next slide, please.

YUYA NAGAI: [inaudible]

YOSHIKAZU KOJIMA: What is difficult in the project was it was first time to use DNSSEC for all project members, except Nagai. So I explained from overview to detail about registrar system, DNS, DNSSEC to the developers of the [web UI] and the backend systems. Especially about DNSSEC, [inaudible] core members of the development team and [inaudible]. Also, I have customer support team manager to understand and make FAQ, etc. Next, I'll explain about transfer policy. Next slide, please.

YUYA NAGAI:

YOSHIKAZU KOJIMA: [inaudible] make sure that how we provide DNSSEC service. If a customer ever to transfer with DNSSEC enabled or not. And we assume [inaudible]. In Japan, JPRS, Japan Registry Service, and DNSSEC Japan community published about the domain transfer with DNSSEC. I took part in this activity. We used this report as reference to decide of our policy. Next slide, please.

YUYA NAGAI: [inaudible]

YOSHIKAZU KOJIMA: In conclusion, we offered to transfer domain without DNSSEC. The reasons are [inaudible]. The first, sometimes the relationship between registrar and DNS provider is not so strong. For example, the registrar recognize the domain is transferred after transfer process is finished.

Second, our customer support receive more support requests. Sometimes, customer forgets that they enabled DNSSEC. I will show about this case in next slide. I have done transfer domain with DNSSEC enabled, and without DNSSEC, DNSSEC disabled, in both cases manually. The transfer domain with DNSSEC enabled was very complicated and troublesome.

If registry, registrar and DNS provider supports the same standard [products] to automate domain transfer with DNSSEC, probably, the situation becomes better. Next slide, please.

YUYA NAGAI: [inaudible]

YOSHIKAZU KOJIMA: Since we provide DNSSEC services, we didn't face serious problem. Here, we show a rare case, some rare case that customer complain. This is case one, the transfer into our service with DNSSEC enabled. This case, customer, then customer complaint that they cannot resolve their DNS records. So the reason was the DS record validating. So the [DNS cache resolver] ignored our DNS response from our authoritative DNS. Then we removed DS record manually in the registry, then problem is solved. Next slide, please.

YUYA NAGAI: [inaudible]

YOSHIKAZU KOJIMA: Second case, the transfer out from our DNS to another service with enabled DNSSEC. Then they asked that they can't resolve the DS record anymore. At that time, they are not our customer anymore but we don't know why they asked us. Anyway, their domain was not under our control, so we cannot do anything for customers. Then we offered to customer that they should ask their current service provider support.

Both case one and case two, they were DNSSEC enabled [at] the process of DNS transfer. Fortunately, we don't have any major problem other than those cases. Next slide, please.

YUYA NAGAI: [inaudible]

YOSHIKAZU KOJIMA: Future [plans and] issues, we don't have feature yet to create DS record on the registries automatically. Recently, we have around 200 domains enabled DNSSEC on our service. So actually, our service has around 2 million domains, so only a few percent is using the DNSSEC. I personally want to provide those services [inaudible] several times for business [customers,] but less customer request means low priority for business. I hope to improve our service better and make customers happy. Thank you.

JACQUES LATOUR: Thank you. I think we need, as an industry, find a good way to transfer signed delegations between registrars and DNS operators or keep working on that. We're a little bit behind time on our schedule, so we're going to go next with Kim Davies on an update with the KSK rollover.

KIM DAVIES: Hi everyone. Thanks for the invitation to present to you today. I was asked to give a condensed version of a presentation I gave recently on some of the challenges we've faced in managing the root zone KSK recently. So this is a condensed version, if you've seen that presentation already, but it's also been updated with some more recent developments. Next slide, please.

First, I'm going to give a brief overview of what normal KSK ceremony operations look like and then go and review some of the challenges we've faced this year in conducting key ceremonies in the normal way. and then I'm going to talk a little bit about the future of how we're considering to hold key ceremonies moving forward. Next slide, please.

A quick primer. The root zone KSK is one of the administrative functions performed by my team as part of the IANA functions, closely related to root zone management. The root zone KSK serves as the trust anchor for DNSSEC.

When we change the KSK, a process known as rollover, it's uniquely complex for us as the root zone KSK because unlike any other zone, you can't just update the delegation signer record and call it a day.

Updating the KSK involves updating trust anchor configurations really around the world in all sorts of validating resolver configurations.

The way we keep the KSK safe is in two geographically distinct facilities, one on the US East Coast and the other on the US West Coast. And the way we store it within these two facilities is we use hardware security modules, and we activate them using an m-of-n scheme. The way that the trust is split in this scheme is we assign responsibility for each of those key shares to trusted community representatives. These are representatives from this community that are distributed all around the world and come to these key ceremonies to exercise their role. Next slide, please.

Some of the security objectives and the way KSK management has been designed, firstly, there's many overlapping layers of security, the idea being that if any one individual security approach is inadequate, there are many other compensating controls there that mean that the security of the system as a whole is not compromised. We really tried to design it in such a fashion that it’s not brittle in the event that any one particular controller is not successful in conducting what it’s trying to accomplish.

The next thing we do is take great efforts to protect the chain of custody of all the critical elements that are used in key ceremonies. This includes not just the KSK itself but all of the materials that are used to act upon it in a key ceremony context. By ensuring that the chain of custody is preserved throughout the lifecycle of all these different materials, it allows us to have confidence to know that the KSK has not been used in a surreptitious or unexpected way.

Another objective of the design is to minimize the risk of collusion. In order to successfully activate the KSK, you need the participation of a large number of people. Some of them are community members, others are different staff members of ICANN who belong in different departments. And really, you need a conspiracy of quite a few people in order to be able to activate the KSK in some fashion, and the risk of collusion is minimized by making sure that pool is quite large.

We have redundancy in our approach. Basically, we have a duplicate of everything, duplicate locations, each one is able to step in for the other if necessary, and even within those locations, we have duplicates, particularly of the HSMs. We have duplicate laptops, duplicate CDs. We have all the materials we use duplicated so that if any one individual piece of equipment fails, we can continue with operations.

We guard against surreptitious entry. Whilst the design of the system is not designed to be Fort Knox in that it is impervious to being entered, what we're really designing against is surreptitious entry, the notion that someone could get in undetected. If we do detect access that is unauthorized, the new can do something about it. But if we do not detect the access, then we have a real problem.

And an open design. All the software that we use in our key ceremonies is public, we've open sourced the software that we've custom built for the ceremony, we use a custom operating system but that operating system build is available to everyone. So that’s all available for inspection. Next slide, please.

So in order to use the KSK, we conduct planned events known as key signing ceremonies. In normal operations, we hold these four times a year. And these are the events where we get those people that are distributed around the world together in order to turn on the HSM and activate it so that we can generate signatures that are used to protect the root zone.

At these ceremonies, these folks get together, and we make sure that whenever the KSK is activated, that it is observed by various different individuals in the room and also remotely. It is audited by a third-party auditor. And the goal here is really to guard against any inadvertent use.

We do these ceremonies in a highly transparent manner. There is a script that we follow. But there is a collegial atmosphere that allows anyone to sort of interject if they feel that there's a concern. If they feel that things are being done improperly or they have a recommendation on how to do things better, that would be taken under consideration, and we might adapt the key ceremony as we go along as a result.

The general purpose of doing it this way rather than in a very closed unwitnessed way is to engender trust in the process. ICANN’s approach has been that DNSSEC will only provide security if the broader community is confident in how the KSK is managed, and we engender that through having this highly transparent approach. Next slide, please.

I mentioned that there is a script that is developed in advance that documents all the detailed step-by-step instructions on how the ceremony is to be performed, and the idea is that that should be something defensible that can be explained to an expert and they can fully understand why it was conducted that way, step by step, there's no extraneous process there, and it is all very transparent. And as I mentioned, when we do these ceremonies, they're streamed online so anyone could follow along in real time. They're also recorded so someone can go back and look at them after the fact. There's also a bunch of people that are actually physically present to witness the ceremony as it goes through. And I mentioned earlier that all of this is available online, the scripts, the code, the videos, all available for inspection by anyone. Next slide, please.

So, what does a good ceremony look like? Essentially, our objective is to do what we need to do to fulfill the objectives of the ceremony and do it without improperly disclosing any of the sensitive materials. There's a variety of different security controls that are part of our third-party audit, and we want to make sure that they're adhered to so that we pass our audit and the community is satisfied that that is done as well.

Now, just because there is a script that we follow doesn’t mean that we have to adhere to the script absolutely. We can deviate from the script if necessary. These are known as exceptions, and they're okay as long as they're properly witnessed and we account for them. Generally speaking, if we run into an issue, just very basic example is maybe there's a typographical error in the script, maybe someone accidentally pulled a power cord out in the middle of a step and that step has to be reperformed, things like that, these can be corrected on the fly relatively straightforwardly. There's no loss of confidence in the system by going off script a little bit to remedy that error. And we still fulfill all the objectives of the ceremony.

And the redundant design of the ceremony that I mentioned helps us with this. If there's some kind of equipment failure, generally speaking, we can then switch to an alternative piece of equipment and still conduct the ceremony as planned.

Ultimately, a good ceremony is something that retains the ongoing community trust in how it’s conducted. Those that were participating, those that were watching, those who review it later, they can all look at what we did and feel satisfied that the KSK was operated upon in a responsible way and they don’t feel there's any significant breach in the security around the KSK. Next slide, please.

So we've been doing these key ceremonies now for ten years. We started in 2010, and that means we've done around 40 of them. And we've really been able to recover from every single issue that we've run into up until recently without any significant challenges. In fact, whenever we hold these ceremonies, all the people that are flying in from the four corners of the world, we always advise them to please be flexible with their travel planning, please be mindful that we may use the next day after the ceremony as a standby day so that if we’re unable to successfully perform a ceremony, we could always come back the following day on the notion that any kind of error we couldn’t overcome in the short-term, we would be able to fix overnight and then be able to do it again the next day.

However, from 2010 to 2019, we never had to use that standby day. We've always been successful in the moment conducting ceremonies. But 2020 has been unique, and in fact, it changed that. So now I'm going to transition into some of our experiences that we've had this year. Next slide, please.

The first thing I'm going to talk about is the first ceremony we did this year, back in February. This was KSK ceremony number 40. Next slide, please. So the key ceremony 40 was scheduled for the 12th of February 2020. The objective here was to do, as we do with every other ceremony, sign three months' worth of signatures for the root zone. In this instance, they would cover the months of April, May and June of this year. And also, we had a secondary objective to decommission an old HSM. Part of the work that we do is constantly upgrading our hardware, retiring it on a cycle, and this was one of those steps where we had some equipment that we no longer used and we needed to decommission it.

We also had some pre-ceremony activity. Some of the maintenance work that we need to do, we don't do in the middle of a ceremony, we do it adjacent to the ceremony. These are performed in a slightly different manner because we don't want to occupy everyone's time during the ceremony with doing minor technical work. But these are held to a high standard, they are audited, there are witnesses, so there are still protections on how they're done. So that was [the plan.] Next slide, please.

So we were doing that pre-ceremony work on the 11th of February. In this instance, it was to upgrade one of the lock assemblies in one of the safes with a newer model. We went to do this ceremony work and the safe would not open. Now, these are electronic locks, they're built to a high specification, and we just could not open it. It does have an electronical display and we were able to ascertain as we tried to open it that the combination was being dialed correctly, but even though the combination was being dialed in correctly, the bolts would not retract in the safe mechanism to allow the door to open.

So ultimately, this meant there was some kind of either electrical or mechanical failure within the lock assembly itself that we weren't able to correct for. Next slide, please.

So the remedy to this problem ultimately was to drill the safe. Now, this is not something that we had ever planned to do. It's something we hypothesized about as one of those disaster recovery scenarios we may have to do in some kind of far fetched extreme scenario, but it's not something we ever really considered would be a likely thing we would need to do. But nonetheless, that's what we had to do. We rapidly brought in expertise to do this, and over the course of 20 hours spread over two days, we did drill into the safe lock assembly. This allowed the bolt to be retracted and allowed the safe to be opened.

Following opening of the safe, the safe was remediated back to its original condition and a new lock assembly was installed, thus returning it back to normal operation.

This was a complicated process. it was also complicated by the fact that the lock is actually designed with a number of anti-defeat mechanisms deliberately designed to make it very hard to drill into the safe. So 20 hours is a long time, and this was done, mind you, with community members and staff members all in the room for those whole 20 hours observing as it was going on. Next slide, please.

One of the things that we were considering during those 20 hours is things like, did someone forget the combination? Did they fumble the mechanisms? These mechanisms are very tricky. And these locks are designed to make it very hard to retry a combination over and over again.

We were obviously considering what is broken. You can't see inside the safe, so there's a lot of hypotheticals being thrown out hour after hour as we went through this process. We were trying to work out obviously how not to break the safe even more than necessary.

Stamina was an issue. You have a bunch of people that are locked in this windowless room hour after hour. Some people stopped being quite collegial at the end of that, and it becomes a challenge. It’s also hard physical labor, what was being done. So this is something that we needed to be mindful of as well.

And lastly, making sure that we maintain the quorum of people there because some people had flights, they needed to get home. Could we manage to do this in time? Next slide, please.

So what happened, while the ceremony was successful in the end, it was held with a four-day delay. We did gain valuable experience, and this will help inform our future disaster recovery plans. The community volunteers and staff alike were very supportive throughout the whole process. And we did take away some learnings from this on how we can improve the way we do these kinds of activities in the future. Next slide, please.

So now I'm going to talk to you about the April 2020 KSP ceremony. This is the subsequent ceremony after the previous one I just discussed. Next slide, please.

So here, the plan was hold a ceremony on the 23rd of April, again signing three months' worth of key material. This time, July, August and September. We also planned here to induct a new HSM, again, part of our refresh cycle for our hardware. And also, we had two of our community volunteers that were wishing to step back from their roles, so we were planning to replace them with two new volunteers that we'd identified from the community.

However, just as we were in the midst of that last ceremony I just discussed in mid-February, we were already becoming aware of the impact of the coronavirus globally, and we immediately shifted our focus to developing contingencies to make sure that we could hold the ceremony even though the pandemic was continuing and the situation was really deteriorating day after day.

So some of the initial work we did here was periodically reevaluating all the participants' ability to travel, continuing to monitor the threat situation, particularly leveraging the resources that ICANN has in its security team, and building out contingency scenarios, constantly talking about what would happen if X happened, what would happen if Y happened? Making sure that we had plans about all the different ways things could pan out.

It's worth noting, before I go any further, that our facilities are designed to let us do ceremonies in a disaster recovery scenario with really a minimum number of people present, even though we desire strongly to have all those people I mentioned earlier present, they are designed to allow us to do ceremonies with a lot less people present if the situation warrants it. However, we'd never really truly defined what those triggering conditions would be for us to go down that path. Next slide, please.

Some of the thoughts that crossed our mind in this planning, [inaudible] people still attend, what if they're sick, what if they can't travel? What if there's government restrictions? What if government won't let us into our facility, let alone travel? What if the owner of the facility won't let us in? These were all very practical issues that we had to consider.

Would we need to drill into the safe deposit boxes within the safe if we couldn't get the TCRs to show up? We do have precedent for this, so we could draw on our expertise in that respect. But that was another thing we considered.

The other question is, let's say we held the ceremony in April. Would we be able to hold it three months later in July, August kind of time frame?

What if the pandemic keeps getting worse and worse? Bear in mind this was back in February. Can we keep our staff self-isolated forever? They need to eventually reintegrate with society. How will things go?

And what happens if we can't hold a ceremony at all? What if it’s just absolutely not possible in any configuration whatsoever? Is the absolute last resort to actually unsign the root zone?

So all this boils down to the fact that we built the system, we designed it so we spread all these people around the world as a security mechanism, but it was actually working against us and it was our worst enemy in this scenario, because the pandemic really meant that we wanted people to be closely proximate to one another to be able to do anything successfully. And how do we do all of this whilst still retaining confidence with everyone? Next slide, please.

So we considered some different ideas. One was holding the ceremony with less than the ideal number of people present, holding the ceremony sooner, the idea is we could sort of beat the peak of this pandemic. Could we post pone it until later? Maybe the pandemic would be very quick and we could do it sooner. Could we move the facility? Could we bring in new community representatives maybe that are a bit more local to the facility? Could we sign things for more than three months at a time? So these are all the things that we considered.

Some of the longer-term ideas that we couldn’t implement for this ceremony but would perhaps solve the issue long-term, reevaluating the locations of our facilities, and also possibly reconfiguring how many TCRs are needed in the normal instance to do the ceremonies.

And lastly, we developed graduated decision process so that there was clear—the triggering conditions that meant we would move on to a certain scenario only when these conditions were satisfied. Next slide, please.

So, what we decided to do was perform the ceremony with a minimum number of people and devise a way of having the ceremony participants broadly participate remotely. We updated our DNSSEC practice statement to clarify the roles and responsibilities in this kind of scenario. We discussed it with the ICANN board of director, the ICANN executive team, and also broadly with the community to obtain buy-in from all these different groups.

We minimized the scope of the ceremony by eliminating all the nonessential work that we were planning to do, really focusing on just on signature generation. And we had four of the seven TCRs transmit their secure credentials in advance to allow four surrogates to perform the ceremony in lieu of them traveling.

We held a ceremony on the same time and date as scheduled, but in El Segundo, not in Culpeper. El Segundo is close to ICANN's headquarters and allowed us to really draw in a lot of staff that were very close to that facility to do it with no unplanned travel.

And what was key also is that we decided to sign nine months' worth of material instead of three. in short, this relieves us of the necessity to do a ceremony now until 2021. The thinking was that should get us clear of the pandemic sufficiently that we can do more planning in advance, we can monitor how things are going, give us some breeding room, essentially, so we didn't have to be pressured into doing a key ceremony right after it in a few months' time. Next slide, please.

So the ceremony was a success. We did it with seven people, and you can see them here. Next slide, please. It was a bare minimum of attendance. Seven was what we deemed we couldn't do it with any fewer due to the controls that we have, and all of them were either staff of PTI or staff of ICANN.

But what we did do is augment the remote participation to make it much more active than it has been historically. All of those that would normally have trusted roles in the ceremony were on a Zoom conference, much like this one, but they were all off mute so they were all able to interject, they were all able to witness it directly. So they had a much more active role than just watching it sort of in a read only fashion.

There was a record attendance on our livestream. I think, off the top of my head, there was some thousands of people watching by the end of it. And most importantly, the ceremony was an absolute success. Next slide, please.

So, that's what we did. Now, what do we need to do in the future? Firstly, general observations, KSK management's highly transparent. We feel that there's high levels of accountability. I mentioned the audit frameworks, I mentioned the role TCRs play, observing, critiquing the process, and all the materials that we use are available for anyone to analyze and review.

We provide thought leadership in this area to others that work in a similar space. We do annual customer satisfaction surveys that are consistently high on how we do these ceremonies. And the events of 2020 have really challenges us with some of the worst-case scenarios that we'd only ever sort of contemplated before but not truly exercised. So tested our ability to be adaptive, allowed us to exercise scenarios that had really only been hypotheticals until then, and stretched us to go the extra mile to maintain high community trust as we went through this process. Next slide, please.

Here are some of the things we think should be active areas of discussion moving forward. One is, where should our key management facilities be, should the locations be rethought, would adding more locations help? Adding more locations makes it more resilient against some of the threats we talked about, but it also increases the attack surface. And they're expensive. And we need to staff them. And if we have more and more facilities and if we continue to rotate through them like we do today, each one lays at rest for a longer and longer period. So these are all things that need to be considered in that kind of discussion.

Is global mobility a thing of the past? Should we rethink this notion that everyone is spread out around the world? This is a key question about the fundamental design of the system. Should we rely more on logical sharing of the trust rather than physical distribution of it? Next slide, please.

Another thing is whether a standby key is warranted. I think one of the things that raced through our mind is if we just couldn't pull together a key ceremony due to localized impacts due to the pandemic affecting somewhere in particular. Standby key might have been another option to consider. So this is just another thing to consider in that debate as well. Next slide, please.

In the midst of everything I just talked about, we were also doing a consultation on how to regularize KSK rollovers. This is work we essentially put on hold for the last months because of all the work we needed to do on these key ceremonies and also sort of our limited capacity due to the coronavirus in general. But nonetheless, we were contemplating that we'd just done a key rollover in 2018 and it was highly successful, how do we keep repeating this on a relatively regular basis to ensure that everything remains agile and everything keeps ticking over?

But we now really need to take all that feedback, that valuable feedback we got from the community, but also augment it with all of the experience we got over the last three months, think about how that new recent experience will impact what we ahead been discussing. Next slide, please.

I think it's worth noting that PTI, who runs the IANA functions, is just on the precipice of approving a strategic plan that will cover from July 2020 through 2024, and it has two specific targeted outcomes that directly speak to this kind of work. Firstly, to adapt to evolving requirements for managing the root KSK, things like evaluating key algorithms, reconfiguring our facilities, and the role of recovery key shares as a viable disaster recovery method.

And then a second one—and this is a relatively new addition—policies and procedures are adopted to ensure successful engagement and future operations despite long-term limits on travel due to COVID-19 and other similar events. So this is really putting a stake in the ground saying these will be key areas of activity for the IANA functions in the next four years, so expect more to come. Next slide, please.

This is my last slide. Constant improvement is part of our DNA. This is an abbreviated version of the slides, but most of with hat I discussed now has evolved over the last ten years. Very little of what we do is the same as we did it in 2010. Constantly renewing our processes, constantly renewing our hardware, always trying to find ways to improve in the way we do things. And if you've listened to this presentation or you're following along and you have new ideas you want us to consider, please do share them with us. We think that the high level of transparency that we use is messy at times, but it really is the best way to ensure success by ensuring that there's a strong positive oversight of the way we do these key ceremonies.

And I think that is the end of my slide deck. Thank you very much.

JACQUES LATOUR: Thank you, Kim. So we're running out of time a little bit here so no questions for your session. Next one is with Dan York with ISOC on Open Standards Everywhere project.

DAN YORK: Hi everyone. Hard to compete with Kim talking about drilling out locks and things like that. Thank you, Kim, for that presentation. That was great to see.

I want to just briefly touch on a new project that we have at the Internet Society because it does include DNSSEC and TLS and pieces around security and parts of that. It's a project that we've called the Open Standards Everywhere project or OSE project. I am the project lead for it.

It began a year ago when we did an audit of our websites and we discovered that we weren't necessarily practicing all of the protocols and all of the pieces that we promote. So one of our IT folks said, "Oh, I can go fix these things up," but a couple weeks later he came back and said, "You know, there's a challenge here. A lot of people want to use the latest standards but they don't know how, or if they do, they don't necessarily understand why, and there's a lot of different sources about individual standards but many of them are very technical, they're not up to date."

So Greg, back then my colleague, said "You need to make this easier for people." So we decided to contribute our contribution here, was this project called the Open Standards Everywhere project, or OSE. What we did was we focused on web servers and we built a set of public web servers that comply with these latest standards, one of which is DNSSEC, providing documentation, promoting this and leading by example.

What we’re focusing on is IPv6, HTTP2, DNSSEC, and then TLS 1.3, HSTS, a number of different parts that come into that. We’re using a test framework and I want to point this out for people. If you have not been to Internet.nl, it’s a great website, a test framework that I would encourage you to go to, test your websites out. It’s run by a consortium of organizations out of the Netherlands and it’s a great place to go and understand what is your website doing, how is it working, what could it be doing better in some ways.

We scoped this project at just the webservers in 2020 and focused on the connections to the server, not focused on the content of the website, so the piece of there, but just focus on what's there in that regard.

We’re also realizing there's really three different types of webservers: some that you host your own, some that are hosted with a webhosting provider, and then others that you use a CDN, content delivery network, in front of it in some way.

So we've built some reference servers throughout there, but perhaps more germane to what I want to talk about here is we’re in the process now of developing documentation that will help people basically provide recipes for the website administrators, the website operators out there to go and be able to go and configure their webservers to use the latest standards such as IPv6, DNSSEC and TLS and pieces like that.

We’re currently developing this documentation on GitHub. We’re using that as—I’d say an experiment partly because we at the Internet Society have not really worked on developing documentation publicly, but already, it’s been great because we've been getting comments, feedback, requests from people.

In the next month or so, we’re going to be moving that to a kind of final state where we’ll then bring it across to our website, translate it into English or into French and Spanish and move it there, but we would invite people to participate, to look at it, to see what's there. We are continuing to work on our won sites, we’re working with our chapters and special interest groups to encourage them to go and also support these different protocols. If you look closely at the slide, you'll see I have a number of sites we still need to finish DNSSEC at.

Interestingly, it has a lot to do with the way we use a content delivery network and we use CNAMEs in pieces in front of things and the correlation between those is tricky to get right in a couple of places. So we’re working through some of that and we’ll be documenting exactly what we do as well.

In the future, this project right now is focused around webservers, and we’re continuing to monitor things like HTTP3, which QUIC is part of it. Some of the website packaging standards. In future years, we intend to go and look at how do we help do the same kind of thing, provide simple how-to recipes for making mail servers more secure. One of the pieces we’re looking at there is how can we help bring in things like DMARC, DKIM and DANE and how do we work with that.

With DNS servers, new things, DOT, DOH, how can we help, again, provide ways to get these configured in this kind of way? Time servers, communication servers.

I would invite you all to go and test your sites with Internet.nl. I think it's a great site that I would encourage you to do, and help spread the word about that site, because the more people who test with it, the more people understand what they need to do to bring their sites into using the latest open Internet standards to make their service more available and more secure.

I would ask you to join us, look at the documentation we have on GitHub. There is documentation there about DNSSEC, about TLS, as well as IPv6 and HTTP2, etc. And if you'd like to know more about this project, you can go to internetsociety.org/ose and you can also contact me directly and I will be around later for some questions if you'd like to ask.

And with that, I know we're kind of crunched for time so I will turn it back over to Jacques and we can move on to the next session. Thank you very much.

JACQUES LATOUR: [Thank you, Dan.] Kathy will share the link to the ICANN website to download the slides from. This is really interesting information and project. Next up is Wes Hardaker from ISI. He's going to talk about carrots and stuff.

WES HARDAKER: Yeah. All right, so I'll keep this brief. I know we're short on time, and that's just fine. I'm going to talk to you today about the local root project, which is a project that I started at ISI a couple years ago and basically, if you think about classic DNS resolution, you have a whole bunch of clients that have to talk to an ISP and the resolver inside an ISP, and that resolver has to navigate the whole DNS tree. We won't go into exactly how DNS works, but I assume everybody knows. You start at the root and work your way down through the various organizations. Next slide, please.

If you're going to ask for example.com, if you have some client that asks for example.com, it’s got to traverse the whole tree. It starts with the root, it says, “Hey, do you have www.example.com?” If you want to see a great skit about this, attend the DNSSEC for beginners workshop, if we ever hold one again.

But the important takeaway is that when the resolver actually gets these answers, it caches the results for a while. So I'm building this cache on the bottom there. you can see that .com and example.com are sort of remembered so it knows where to go ask next time. Next slide, please.

So if you need to go to ICANN for example, you go to www.icann.org, it doesn’t know anything about any of that so it has to start at the root all over again, ask for where’s .org, where’s ICANN.org, and build up this whole tree. Next slide, please.

If you need, on the other hand, to go ask for exam.com, well, the nice thing is it already knows where .com is. It doesn’t have to go back to the root to ask for .com because before, somebody had asked for example.com so it knew exactly where .com was. So it stats off and has to go ask .com where example.com is, because it didn't know that. But it’s able to skip some things. And this is highly beneficial. Next slide, please.

So the real question is, what if we could pre-cache everything? What if we actually knew all of the answers? The best way to be fast is to always know the answers ahead of time. You don’t have to go query somebody else. Next slide, please.

That’s exactly what local root does. It’s a pseudo cache where it is pre-caching a lot of the information within the root zone or other zones, as we’ll see in a minute. So you already know where .com is, you already know where .org is, you already know where .pr is, you already know where .horses is and you can immediately jump to the right place and never talk to the root, which is why there's a big red X. Basically, you never have to talk to the root again except to pull the root zone. Next slide, please.

So this extends RFC7706 which was just replaced this week, actually, by RFC8806, on Friday in fact. So brand new information, it’s just an update to the same IETF specification about this whole concept. We add a few things to this in the local root project. We add notifications. 8806 doesn’t handle notifications to let your resolver know, “Hey, now is the time to go pull new data.” They’ll get it anyway eventually, but ours is a little bit faster and you get notifications out of it. And then we also do a secured transfer using TSIG keys, which has some benefits I'll get into in a second. Next slide, please.

So, why use local root? As I already talked about, you get this sort of pseudo caching of the root and some other zones. We’ll talk more about that in a second. And it removes the need to contact them frequently. And that really boils down to two things. One, you get faster DNS lookups of the first TLD and other related lookups. So you already know the answer, you don't have to go reach out to the Internet to find the answer. And you also get faster NX domain results. So negative answers, in other words, when you ask for something that doesn't exist, if you ask for .horses with two H, that doesn't actually exist as a TLD. It takes a while for you to ask and be told, "No, that doesn't exist."

Looking in the recent DITL data, which is a day in the life of the Internet collection done by DNS OARC for one server, there were 6.7 billion requests in one day to the root zone and only 1.34 were actually valid. The rest of those were actually nonexistent domains. So it's actually helpful to get faster nonexistent domains just as much as it is helpful to get faster real domains.

So, what else can you do? Well, this is sort of built as a test platform, people can sort of do what they want with it as well. the notifications might actually be beneficial to people doing research on how often does the root zone change or things like that. So we can tie all that together. Next slide, please.

So there's been a bunch of recent improvements. As I said, I started this two years ago or so, and thanks to some sponsorship by ICANN and some push forward within ISI, we really rolled this out into a production system now whereas before, it was sort of an experimental effort.

First off, we added IPv6 support. This was one of the earliest requests I had. I only had a v4 enabled server at the time. Now we also have three upstream servers. Before, we only had one server that your resolver could talk to. Now we have three servers and they're all running both IPv4 and IPv6. And the end result is that your resolver can get this local root precache installed by downloading stuff from the upstream local root servers. And then the rest of your DNS requests, of course, still have to go out to the Internet as indicated by the bottom half of the diagram. Next slide, please.

So a bunch of recent improvements have just been pushed out. And interestingly enough, the COVID-19 virus actually affected this. I hoped to have this done much earlier. But one of the things we did was actually redeploy some new infrastructure and did some [IP] renumbering and things like that. And the delays for people actually being able to go out and pull hardware and install servers and stuff did slow us down a little bit, unfortunately. But we're up and running now.

So we added configuration support for BIND, unbound and NSD as resolver types that automatically generates configuration for you for those three resolvers. We have multiple zones supported, so not only do you get the root zone, you get .arpa, root-servers.net, and DNSSEC-tools.org as well, and we'll come back to that in a minute.

There's user preferences now so that you can actually control a little bit about your account and what sort of notifications you get, and then we send notifications over e-mail for things like announcements or required configuration changes and things like that so you can be kept up to date.

And then the whole system has been moved inside our production service monitoring so that we are keeping very close eye on local root to make sure that it's always up and always available and we get notified and woken up in the middle of the night if it's not true. And then a lot of UI and documentation improvements. Next slide, please.

So this is what the home page looks like. It's at localroot.isi.edu. There'll be a URL at the end of the slides as well. You'll notice in the upper right there's a register and login button and there's some news that you can see that was about the major update. Next slide, please.

When you register and you log in, you get this sort of screen that shows you your list of servers, you can add a server, you can create configuration and things like that. Note the "get config" button on the right-hand side really is—this is the most powerful aspect. Once you register your server, it just gives you everything you need to stand up a new resolver. Next slide, please.

So this is the configuration generation option. So when you click on that "get config" button, the three things in red there are all brand new. It asks you what type of resolver you want to generate configuration for, on the right-hand side, it asks you what zones you want to mirror. By default, DNSSEC-tools.org is off but the rest of them are on, but you can pick and choose between what you want, and then it'll help you configure your configuration for what addresses you want to listen to. Do you want to just listen to the loopback address, which is what 8806 talks about, or do you want to listen to some other internal addresses as well for your clients? Next slide, please.

So this is actually what the configuration output looks like. You'll see that this is for unbound, and it gives everything that you need—there is the upstream servers listed in red are the three upstream servers that local root provides for both v4 and v6, and then there's a bunch of backup servers that the rest of the roots will supply you data swell. Next slide, please.

This is the account preferences. It’s allowing you to check and uncheck whether you want e-mail on various types of things like local root news or configuration changes. There’ll be more on this slide in a minute. I'll get to that in a second. Next slide, please. So really, there's a few things. We've accomplished a lot in the last six months or so, but our goal is to push forward with it now that all this base infrastructure is in place.

One of the things I want to do is if your resolver—we haven't seen you ask for a data update in a while, we can actually go out and tell you it’s been too long, you might want to check on it. So, why is this important for DNSSEC? Well, as you know, DNSSEC requires that signatures be updated on a regular basis. So if you don’t have the most recent data, eventually, you might fall into a DNSSEC invalidation period where your resolver no longer has the up-to-date data.

So we want to be able to warn you when we haven't seen you transfer data in a while that you really need to pull it quickly and get that configuration fixed.

We also want to support some other small to medium zones. So besides just the four zones that I mentioned, I would love to hear from other zone owners that have data that you think is widely needed and generally small enough to warrant this type of service, maybe other small TLDs or something like this so that ISPs can enable this service for your zone as well. And if you have the interest in doing that, please contact me. My address is down there at the bottom.

We also push out some other things, like a RUST API and group accounts and things like that. Those are coming soon as well. If you have any feature requests or have feedback on this service, please e-mail me. Again, the address for it is localroot.isi.edu which I thought was on this last slide, but it's not in this copy. It is on the front slide if you want to jump back to the front, Kathy. But other than that, I think I'm done and we've probably saved some more time. It's not there either. Darn it.

JACQUES LATOUR: Thank you, Wes. Now I will turn over to Steve for the DNSSEC provisioning panel.

STEVE CROCKER: Thank you, Jacques. I know we're running behind. I have a full panel here. So I will just kick it off as quickly as possible. There we go. Share. And now everybody should be seeing this. We'll set this up as a slide show. Okay, this is the second in a series of panels on automating DNSSEC provisioning. Today, we're going to focus on two parts.

Shumon Huque and I have put this together. Shumon has been doing some fantastic work. We're focusing on two aspects. One is on the automation of the updates of the DS records and the other is coordination across multiple zone providers, multiple DNS operators, and how to coordinate all of that.

A particularly interesting and important subcase is that this also covers how to transition from one DNS operator to another. So the problem that you heard in earlier presentation today is a piece of what's being addressed here. That’s really addressing a more general problem.

So today’s panel, we have people from registries, registrars and DNS providers on the cross-signing, the key piece of work is a new document from Shumon and company, multi-signer DNSSEC models, it is an RFC in progress, I think. Shumon will talk about that quite a bit.

The other aspect is the topic that we took up last time and we’re continuing on, automating DS updates from third party providers. And the questions are, where do we stand on all this, what are the next steps for automation, and what are the impediments.

Our panelists today, I will stop talking momentarily, turn things over to Shumon, and then Paul Ebersman from Neustar, Brian Dickson from GoDaddy, Jothan Frakes from PLISK, Jaromir Talir from CZNIC, and Oli Shacher from SWITCH. And as you can see, these are representing multiple different functions here. Brian is straddling between being both a registrar and a DNS provider in this.

Very briefly, the problem of conveying a DS key up to the registry can be done in roughly one of two ways. One is either to push it upward, which is symbolized in blue, or to pull it by having either the registry or the registrar or the registrant go down to the DNS provider and fetch it from the zone there.

The dotted lines indicate that those are theoretically possible but no known instances of them working. The only combination that seems to be working these days is the pull version in which the CDS or the CDNS keys are published in the zone, and the registries pull that up. There's an RFC8078 that governs that, and that's where we stand.

With respect to cross signing between multiple DNS providers, there's again a couple of different ways one could imagine that it would work. The critical thing is communicating the ZSKs across from one to the other so that they can cross sign each other's zones and when they update the keys, they have to communicate them again.

You can either have this controlled by the registrant or some software or a service running on behalf of the registrant or one might imagine introducing the operators to each other and having them self-organize. There are no known examples of the latter, so we'll only focus on the first which is having the registrant or a service running on behalf of the registrant doing the coordination and interacting with each of the zones. As I say, this is the most likely and the self-organizing one doesn't seem to be very likely.

We have a mailing list, DNSSEC-provisioning@shinkuro.com which is focused entirely on these class of problems, and with that, I will turn things over to Shumon. I will control the slides, and we will move smoothly from one to the other.

SHUMON HUQUE: Thank you, Steve. Hello, folks. It’s good to be here, and today, I will give you an overview of multi-signer DNSSEC. Next slide, please. So, as you might know, there's an increased trend in the industry towards using multiple DNS providers or operators for obvious reasons of increased redundancy and survivability. Now, the two most common ways to deploy DNSSEC in such a configuration are, one, the traditional zone transfer model, and two, what I'm going to call the provider API model.

The zone transfer model is actually well understood and it works fine with DNSSEC. Next slide, please. But organizations often want to use a set of non-standardized dynamic features that cannot be supported by zone transfer.

Or another reason, they may not even want to run a backend primary signing server and want to delegate that task to the operators that they’ve contracted with to run the zone for them. So for that, we need to make DNSSEC work with the other model where each provider independently signs the zone data with their own DNSSEC keys, and that'll be the focus of my remarks today. Next slide, please.

You may be wondering what these dynamic features are, so I'll elaborate very briefly. These are things like GSLB, global server load balancing, probing and failover records, weighted responses or even custom programmed dynamic responses. So they're often querier specific or dependent on inspecting some sort of dynamic state in the network, so the answer and signature really have to be determined at the authoritative servers themselves, or at the time of the query. And this necessary means the provider needs to sing the zone with their own keys. Next slide, please.

That brings me to the topic of the multi-signer models. Next slide, please. Our situation is that the zone owner uses provider-specific APIs to update content identically at each of a number of providers and each provider then signs the zone data with their own keys.

For this configuration to actually work, a new set of key management mechanisms had to be defined and all of this is described in the multi-signer model specification which I've provided a link for here, so as you heard Steve mention, this pack has been approved by the IETF and will be published as an RFC quite soon, I think. Next slide, please.

The main requirement in the multi-signer models is that we need to manage the contents of the DNS key and the DS sets in such a way that it's always possible for a resolver to validate the answer no matter which provider it came from. And this requirement is satisfied by having each provider import the public portion of the zone signing key of every other provider into their DNS key sets. Next slide, please.

Two models have been developed. In the first one, the zone owner holds a common key signing key, so that's the gray key at the top in the middle of the diagram, and each provider has their own ZSK. So there's a blue provider on the left and a red provider on the right of this diagram. Next slide, please.

And what happens is the zone owner uses each provider's API to obtain their respective ZSKs, builds and signs a resulting DNS key set, and then pushes the DNS key set back to each provider. And the DS record set in the parent references just the common zone owner KSK. Once we bootstrap the providers in this way, everything is ready to go and DNSSEC basically just works. Next slide, please.

So the second model is where each provider has their own KSK and ZSK sets, and the zone owner's task here is essentially to coordinate the cross sharing of the ZSKs between the providers. Next slide, please.

So, ZSK from provider A has to go to provider B, and next slide, you'll see the vice versa, provider B ZSK is going to the other side too. At the moment, this requires the zone owner to coordinate the cross sharing, but I think there's an opportunity which Steve hinted at of devising a protocol to automate this directly between the providers themselves. I've started some preliminary discussions on this topic with a few others, so we'll see where that goes later, but it doesn't exist today, basically. Next slide, please.

And then finally, in this configuration, the DS set in the parent zone needs to reference each provider's KSK, so the blue KSK and the red KSK has to appear in the DS set. And as with the first model, once this configuration boot strapping is in place, DNSSEC just works. There are some additional operational tasks when you're doing things like key rollovers, but those are just details which I'm going to omit for now in the interest of time. Next slide, please.

There are commonly used software toolkits out there that help a zone owner kind of consistently manage zone content across multiple providers and ensure that everything is always in sync. So you might have heard, or maybe you've already used tools like OctoDNS, Denominator, Terraform. There are a few others. And I think enhancements could likely be made to each of them to support these new multi-signer key management mechanisms. So that's an area of investigation for us. Next slide, please.

Let's talk a little bit about implementation and deployment. The protocol is still quite new, but there's been a fair amount of testing in hackathons and lab environments and a set of test zones. NS1 has a production implementation already. We were actually originally expecting one of our colleagues, [Jan Chelak] from NS1 to participate in this panel, but he just became a father and rightly had to excuse himself this time. But there are at least two other major DNS vendors I know that are currently working on implementations. Next slide, please.

I'm not going to go over this diagram in detail, but I just plucked it out of NS1's own documentation. This is essentially their API and the extensions they've made to support the multi signer mechanisms. Let's move on to what I think is my final slide, if you could. One more. Which is the topic of automating the provisioning and management of delegation signer records.

Now, model one is probably a little bit easier in the sense that the zone owner still holds the common KSK, so they are automatically primarily responsible for updating the DNS contents. CDS and CDNSKEY can work fine in this configuration if the relevant TLDs and/or registrars support them, which I know is not often the case today, extending multi-provider toolkits to talk to registrars could probably help in a big way.

So one of our panelists today you'll hear from a little bit later, Jothan, has informed me recently that Hexonet for example is a [SaaS] backend for several hundred registrars. So developing plugins to talk to Hexonet and other key registrar systems could be a first step and could prove to be very useful.

In model two where the DNS providers hold the KSKs themselves, a little bit more coordination with the zone owner is needed. Again, CDS and CDNSKEY can be made to work here with that additional coordination, but it might be worth thinking about alternative solutions also, such as the use of registrar mediated protocols. In the past, DomainConnect has been mentioned as a possibility, but I've also learned recently from Brian, one of our other panelists today, that it may not be the case. So maybe we can ask Brian to elaborate there. But there may be other systems that are available or could be developed that could provide delegated authorization to DNS operators to kind of surgically update only the DS record piece.

And then lastly, the idea of formally designating DNS operators in the RRR system as kind of first class citizens has come up now and again. I know Steve seems to be a bit pessimistic about that possibility, but I'm mentioning it in case it prompts some further discussion amongst the panelists. I think that was my last slide. Did I have anything more? Yeah, that’s it, so I'll stop and turn it back to Steve. Thank you.

STEVE CROCKER: Thank you. Very nicely done. Paul Ebersman, your turn.

PAUL EBERSMAN: Thank you. Next slide, please. So Neustar has been working on the multi-signer draft as well. There were a number of things that we didn't have originally that we needed to develop, and in particular, we're one of those folks that have the advanced features that folks [like and that] do not fit into standard DNS wire protocol. So we had to do a number of things.

The biggest things, starting off with being able to actually take in keys. Our signing stuff was all offline static signing, so we had no way, at the time, of gating private keying material so we had to put that into our code.

The next big steps which we wanted for a number of other reasons was we started changing the algorithm that we used for DNSSEC signing. ECDSA not only has very good size features, for the same cryptographic strength requires many less bits on the wire, which was attractive. But the other thing that it has is it is cryptographically relatively cheap on CPU to generate. So in the past, it was just too slow to be able to sign things in real time. And with ECDSA, we were actually able to have a fairly negligible effect so that every one of our authoritative servers could actually have a private key and sign records as they were generated. That allowed us to be able to do all of the rules and filtering and ACLs and all of the other things that we do for GeoIP and load balancing and other features, finally decide what is the record that we’re actually going to put in the response, and then sign it and send it out.

So, that really was probably the biggest amount of work. The other thing then of course is that once we have internal support for generating or importing keys and being able to export keys, we need them accessible via an API. Most of the folks that want to do multi-provider DNS and all the rest have also dived into all of the dev ops that we all do and have started using containers and VMs and various other things. So most of our larger customers, if they use our web UI at all, are doing it or a temporary tweak. The vast majority of what they do needs to actually be accessible via that API.

We also looked at some other things. CDS is certainly something that we would like to be able to do. There's currently not universal support for that. And currently, our authoritative product has us as a DNS operator only. We are not a credentialed registry/registrar operator with ICANN for that piece. We have that business but we’re actually in the midst of divesting that. So we can't get into those, and as Shumon was saying, the idea that pure DNS operators where there's a registry, a registrar and a DNS operator and none of them are the same entity, and having DNS operators being a full class citizen within the ICANN community has some challenges. And probably, the biggest one is that with contracted parties, the generic TLDs, we could potentially have that become part of the contract just like DNSSEC support has become part of the contract.

Unfortunately, ccTLDs are their own little breed and they can do whatever they choose. Some ccTLDs are at the absolute cutting edge of feature functionality and security. Others of them are lagging dramatically behind. So that really hit us as a problem. So we're still looking at what we can do. At the moment, the reality is that API support and the customer having a fair amount of sophistication and control is still what happens. Next slide, please.

So right now, it's more likely that we'll have more folks asking us for the second model in the draft where there are separate KSKs and ZSKs for each side. And that has some advantages in that we can go through and we don't have to have algorithmic consistency across all of the chain, but it also means that key rollovers for KSK rollovers are going to be interesting.

One of the things we're looking at is what happens when you have as many as four sets of KSKs and four sets of DSes as both providers are potentially rolling. Response sizes are icky. There are also bits of fun like DS records with some TLDs have up to a 48-hour TTL. So coordinating all of those changes in the window during which you have to have all of those things in both sets of keys available so that you don't have things with cache inconsistencies causing DNSSEC validation. That window is pretty wide. When you tell somebody that they can't do this in a weekend cleanly, that's a problem.

And then at the moment, yeah, we're still fighting with all of the various ways in which registries and registrars all do and don't support DS, and some even still prefer DNS key. So we have had a lot of issues with figuring out which to support or how we support all of them, and we've spent probably almost as much time on documentation and assistance for customers in detailed check lists on how to do all of these things as we have writing code. And I think that covers my piece of it, Steve. Turn it back over to you.

STEVE CROCKER: Thank you very much. We move on to Brian Dixon from GoDaddy.

BRIAN DIXON: Hi. We're going to talk a little bit about what we do and what we're interested in helping on the protocol development and a lot of other related things for both us as a DNS operator and as a registrar. Next slide, please.

So what we currently offer for our registrar customers as an integrated part of our product is where we're doing both registrar and DNS provider. When somebody asks us to do DNSSEC for them, we manage everything. It's a one-click, one-button thing. We submit the DS record through EPP and all of the support, the maintenance related to that happens in the background with automatic triggers for rolling the ZSKs and less frequently the KSKs, and updates to the registry.

We always do CDS and CDNSKEY on the offhand chance that registries decide to support that. We know that there are some that do that but everybody. And we're looking at what the roadblocks are, what the obstacles are for much wider deployment in the industry as well as from our customer base for getting DNSSEC rolled out on a much bigger basis. So that’s where some of the things that I've been looking at behind the scenes talking to folks about different ideas for handling some of the cross signing as well as potential ways of making the polling side of CDS and CDNSKEY more scalable.

So those are things that are going to be areas of investigation and collaboration and proof of concept-type work. Nothing really to be discussed beyond that right now other than stuff we’re going to work on. And then we do offer zone transfers in and out with a limited set of combinations, and we’re going to be looking at how to expand that to support the possible case where we transfer zones in and then sign them. Right now, transfers in, we only support transfer in of already signed zones, which is not ideal for the cases such as the special case of case one that Shumon was talking about. As well as the issues about cross signing.

We don’t do sign on the fly, so a lot of the restrictions that require the multiple signer don’t really apply to us, but we want to play well with others and we also want to make it possible to increase the reach of registrars and registries that are able to do DNSSEC. Next slide, please.

I think I mentioned most of this. We just want to remind folks that DS is really only strictly needed to activate if registries are able to support CDS and CDNSKEY. Once the initial DS is there, it should be possible to automate through the CDS and CDNSKEY and it maintains the security model since you have to have the key in order to do an update to the DS records and you have to prove that you own the key. So the only bottleneck is the initial configuration of DS records. And we’re looking at some potential ways of leveraging things like DNS over TLS to authority as potential way of having a registry talk to a DNS provider to get in a secure fashion a DS record from the DNS provider. This is something that's way before even proof of concept, but it's a concept for getting the DNS operator in as a first-class citizen.

The real obstacle I see is that there are a lot of registries and there's a lot of registrars, and even if there is an EPP extension for DS and DNS key, implementing that and making it widely available, especially from a registrar side, generally is going to require either a UI or an API to get the data in from a DNS provider. That's not something that's getting traction quickly enough in our opinion, or my opinion anyway. So looking at getting the DNS providers in as a first-class citizen is a way to get more traction to kind of put more rocket boosters on the whole thing and get it moving faster. We see the benefits of DNSSEC as being very substantial, and we're in a strong position to help make a lot of things happen, but we want to keep it neutral. Next slide, please.

So there's some ideas I've been bouncing around with some other folks about ways of doing that without actually requiring a literal API by using in-band DNS queries as a mechanism for having the zone owner get information from the providers and then submitting that through their zone transfers to the providers, and that creates a closed loop system which is still secure, but avoids the need to have a lot of APIs defined. This simplifies the whole thing, especially for us. We're happy with stuff working over the DNS in band. So I think that would work fairly well. Next slide, please.

And really, the one issue we see as just being a problem is that right now, there's no simple easy button, that if you're not a registrar customer who’s using the registrar as your DNS provider, it’s not smooth or easy to do. But it would be nice if it was a lot easier or more scalable, and that’s the idea of having the DNS provider being a first-class citizen, and then potentially using DANE, TLS A records for establishing the TLS certificate that a registry could talk to a DNS provider and then have a secure channel over which you can pull the DS record once it knows that it’s supposed to. And that’s one of the areas where we’re just exploring the capabilities and seeing what that allows and how to make that more scalable. The other idea is having a per-DNS provider inventory of zones that they manage that are going DNSSEC-secure so that it alleviates the need for having the EPP transaction occur.

The registries already have the NS records and that pretty much, once you have DANE TLSA, that’s all you need. You know who the nameserver names are, they're authoritative within the database on the registry. So the registry has everything it needs to find and talk to the DNS provider. So it’s just a question of populating the right kinds of information and making a more formal situation and training the registries how to do this. Next slide, please.

STEVE CROCKER: Thank you, and now Jothan Frakes from PLISK.

JOTHAN FRAKES: Hi everyone. Thanks. What a privilege and honor, and thanks for staying up late, for those of you in the east coast of the US. Great to see you all. I can go very quickly through this. I think we've covered a lot of the ground here very well.

In talking within the context of the specific sort of path A where we would be having an interaction whereby aspects of domain names related to the DNSSEC configuration can be done through an API. Next slide, Steve.

I wanted to cover that a little bit and just keep it very high level. And just for context, just as though you're walking through a shopping mall or something, if you'll note that red A, that's where we're at. That's what I'm focusing on a little bit here and wanted to just talk briefly about registrars and some paths that do exist that could be a way forward.

But I did want to give the context, and [inaudible] in the header, that this is for technical power users. This is for people who are probably higher altitude than most consumers.

And as Brian touched a little bit about GoDaddy, and my registrar [inaudible] does this, as well as most registrars you'll find in the space, folks who come into configure a domain name are not typically power users, so we have to make things incredibly simple for them and I would put DNSSEC in simple. They're not synonyms. They're antonyms. Unless you really get into it and work with it at the scale that many of us do in this community.

So, in talking with Shumon about the approach that would need to be taken in order to do multi-sig or to make it easier for the person who's managing the DNS, whether it's a third-party provider, an enterprise IT department, even the end user, to interface with something through the API or just through a website is somewhat complex but there are resources and tools available.

The diagram you have here on the screen shows a couple of sets of circles. In the top four circles, we're talking about what are called the registration path and then the bottom circles are the resolution path and they're separated but you'll notice there are some important tethers between those.

The way that all of this DNSSEC magic works is because those have records that are operated and maintained in harmony. And the ability for somebody in the type A or the API or abstracted integration could manage DNSSEC host records or DNS records through an API is possible with a number of different registrars.

There is a reticent to the second path with what's being called the first class citizen path, simply because it introduces new parties to things that have been very, very well structured over a couple of decades and there are some impacts. There's not a reticence necessarily, but there is a concern because everything is architected around the registry and registrar directly communicating.

So, getting that enterprise IT, the third-party DNS provider or someone in to be able to manage or work with DNSSEC records, to be able to update the IP addresses of host record, to add [glue] records or manage those, to change the name servers for domain name, registrars do make that available directly through websites. But they often sell products or services that could be disrupted for customers, and so they want to use caution and care.

The API method is a fantastic way to get there, to bridge these, and so we're talking about a very, very narrow, highly adept and technically sophisticated power user being able to access a variety of different APIs to accomplish what they want to accomplish.

Now, there are hundreds of registrars. I had mentioned to Shuman the technology provider I work with is called [Hexon at oneAPI] but there is many who offer APIs where a user or a system could interact with those systems.

I've taken, and on the next slide, I've collected many of those. I did a callout to the Registrar Stakeholder Group and I mentioned—was invited to speak on this panel and I had the opportunity to share some of the resources that are available with DNSSEC APIs and I've made them available here on this chart, on this slide, as links. So, if you download this presentation, you'll be able to access those directly and be able to look at them further to see if they suit your needs.

As with many technical solutions and the diversity of markets and languages and technical approaches, each provider is pretty distinct in how they approach their API but the APIs are there and available and there is more opportunity to research these available at these various links, email addresses or other manners.

I meant to go very quickly as to perhaps collect some of the [time] and I'll leave any time I might have had in addition to this towards questions or that we can maybe narrow the gap on our time, so that our other important speakers have an opportunity to talk. Thank you very much.

STEVE CROCKER: Thank you. And indeed all the presentations will be available, if they're not already, on the website. All right. We now move from late night on the East Coast of the US to early morning in Europe. Jaromir, you're up.

JAROMIR TALIR: Good morning. Good morning from Prague and thank you Steve for inviting me to this panel. I would like to give you a quick update about the status of DNSSEC automation in dot-cz, so please, next slide.

We actually started three years ago. We have this June we have three years anniversary, so we have some experience for a time. And what we actually do is that we do daily scanning of our hole dot-cz zone and we are scanning for a CDNSKEY just from the single location at the moment.

And we are waiting for seven days, checking whether there is not a change in the receiving data, and if this is the case, we create a DS record for the scanned domain.

We also notify the registrant and the technical contacts about this change and we update registrars via EPP mechanisms, poll messages, unfortunately in a not standard way.

The whole tools, the set of tools is part of our open source registry called FRED, which is currently used by nine countries, and one of these countries started also to use this tool to run CDNSKEY scanning. It's Costa Rica. And some others are testing this. So, I hope that in a short future, there will be much more countries, many more countries. Next slide, please.

We already have roughly 15,000 domains with automation enabled, which is sort of like 5-6% of the total of the dot-cz signed domains. At the beginning, there was a slow start, but you can see we have two spikes in earlier years. And surprisingly, these two spikes are from registrars that already have the EPP access. So even though this tool was, at the beginning, meant for the DNS operators that don't have EPP access as the way how to update DS records to the registry, even the registrars find this interesting.

And it actually makes sense, because if they use the tools, like [RNO] DNS which provides the full [KSK] automation, it's much easier for them to just switch to [KSK] automation on and they don't need to care about anything else. They don't need to synchronize the updates of the zone with the EPP because they can rely on the fact that we are scanning.

So, maybe in the future more and more registrars will switch to this approach as well. Next slide.

So, after the three years of experience, we also think about rebuilding our infrastructure for CDNSKEY scanning, and much of this work is being done as part of the master thesis of my colleague Marina Shchalava. Some features of the new tool is, of course, we would like to add scanning from multiple points, not just single, and this allows us to reduce the seven days maybe to less, to just a few.

If you would like to a little bit [inaudible] tools from our registry system, so it would be useful for other registries as well. And there are some other features that you can actually read in the master thesis because it has been successfully [inaudible] last week, so it's already public and I believe that the output of the work of my colleague will be available late during this year. Next slide.

Okay. That's the update from dot-cz. Thank you, Steve.

STEVE CROCKER: Very well done. Thank you very much. And we move to [Ali Shocker]. Did I say it right?

[ALI SHOCKER]: Yeah, that's correct. Can you hear me?

STEVE CROCKER: Yeah. This is you.

[ALI SHOCKER]: Yeah. Perfect. Okay. So, thank you and good morning from Turk. So, we can go to the next slide already I think, please.

We have a very similar set of tools [inaudible] which is no wonder because we just often can [inaudible]. So, we have also a scanner that runs every day. It runs on three different locations, so we already have that in place that we scan for multiple locations and those scanners [inaudible] outputs to the registry application.

It will compare all the results, perform security checks, and then activate the DS records after three days of consistent publication. So, that's one of the differences to CZ. We only have three days for activation time. We think that's, as [inaudible] already said, that the scanning from multiple locations allows us to [inaudible] harder and then there is no additional benefit in waiting any longer.

We scan for CDS instead of CDNSKEY. That's also a difference to CZ. This is mainly to be consistent with our EPP where we historically have only processed DS records from registrars. And it seems inconsistent to do the other thing while we're [over CDNSKEY].

Another difference is we don't send any mails to the registrants because we don't actually have a direct contractual relationship with them and they think this would cause mostly confusion if you send the domain holders the mail while we are updating your DNSSEC configuration. But we do provide a website where everyone can answer their domain name and check the current state of their domain, so they will see the activation date, when we will activate the change, or if there are any errors.

Once we activate the change, we will then notify the registrars over EPP there is now, there is a standard extension for this kind of notification where we can tell the registrars that you have changed something on the registry side, so if they have some kind of local cash state for the domain, they can refresh [inaudible]. Next slide, please.

So, the adoption. When we started this project in 2018, we checked the number of published CDS records and there were only 640 dot-ch domains. All of them were hosted by CloudFlare at the time, but in the last two years, actually several hosters which are mostly using [knot or PowerDNS], they have started publishing CTS records for all their domains and this is posting the number to over 13,000 now.

About 12,000 things have actually been bootstrapped over CTS, so we're using this process and the rest already had their DS in the parent and there could be that maybe the DS algorithm, the hash algorithm changed or something, but those were not on [inaudible].

And as you can see, CDS is currently mostly used for bootstrapping and we only have seen a few dozen actual key rollovers or the [inaudible]. This is also because there is not much software around which actually supports the [inaudible] over CDS. I think [knot] is the only one that does this in a very easy way, currently. Next slide, please.

There is some work in process. The system has been running smoothly for over a year now, but now we are planning two changes. So, the first is an additional security check we did not have so far. This makes sure that the CDS [inaudible] that is actually signed by KSK and not just [ZSK]. This was recommended by Tony Finch in his article about the effects on [inaudible] and their consequences for DNSSEC. And yeah, we have already implemented this test but we found a few hosters with several hundred domains published in CDS records which are using old

[knot] version as their signer and this old [knot] version signs the CDS actually with the [ZSK].

So we are coordinating with them currently, that they can either update their software or be aware that their CDS signals will no longer be accepted if they don't do that. So we want to coordinate with them before we activate the change and then enforce this.

The second change we're working on is support for EdDSA algorithms 15 and 16. We currently only accept them over EPP, but our CDS processing change includes the resolvers and software libraries which do not support, validate [inaudible] algorithms. So we have to operate this.

I think that's the short version of our CDS story. Thank you.

STEVE CROCKER: Thank you. So, I think … What's happening here? I think that's the end of it, right? So, that brings us to the end of our panel. Fantastic work by all of you. Thank you very much for all of the work that you put into this. This is presented very quickly but a lot of preparation, so thank you, Ali and Jaromir, Jothan, Brian, Paul, and Shumon. Back to you, Kathy, I think. Who's in charge?

JACQUES LATOUR: I think we have a little bit of time for questions.

STEVE CROCKER: Oh. Jacques. Yes.

JACQUES LATOUR: Yes. We've got some time for questions. You can post them in the chat tool. If not …

[KATHY SCHNITT]: Jacques, I do have one from earlier if you want me to just go ahead and read it. This is from actually when we started today.

"Question. Not necessarily for the current subject or topic. Every Internet user is concerned about his or her DNS security, whether or not he or she is familiar with the term DNS. The task of implementing DNSSEC across the Internet would be more complete is there is far greater user awareness and involvement of DNSSEC. Is there a possibility that SSAC develops a user-level security suite for opensource and proprietary operating systems that includes rudimentary and simple fixes or upgrades to enable by default user-level secure settings? DoH, for example. And resolver, cache settings, browser integrity, check settings, etc. Such a fix can be programmed by the user, by DNSSEC dot-x file, dot-dev file, or a BSD port that the user can install by a mouse click. As an alternative to SSAC initializing an exercise to develop this suite, SSAC could liaise with an opensource proprietary developer community in this direction. Such an initiative could include related objectives such as IPv6 and UA readiness." End question.

JACQUES LATOUR: I don't know where to start.

DAN YORK: Actually, there's been a number of those kind of projects that have gone on. I see Steve smiling because this was a large part of what the DNSSEC deployment initiative focused on for a good bit of time. At the Internet Society, we focused on it with our Deploy360 program for a bit, continuing out with the Open Standards Everywhere Project in a different form.

I think the challenge of getting protocols, getting things like this deployed are challenges, especially in ones where the business case is not immediately clear and that's I think the challenge that we have with where we are with some of this.

As to whether it's something SSAC can do, that would be something that SSAC would have to take up and talk about a bit more.

But thank you for the feedback I think we should say, certainly.

STEVE CROCKER: I don't speak for SSAC anymore but I would think this falls more in your bailiwick, Dan, with your opensource [inaudible] everywhere.

DAN YORK: Yeah. That was actually what I was thinking. There's other … It seems outside of where SSAC goes, it's more to organizations like the Internet

Society and other organizations who are out there working to promote this.

I would say to the person who submitted it, if you have contacts within the opensource or developer community that you think would be useful ones to help encourage them to do that, I would say we already have people on this call, such as Jaromir I see from CZNIC, and some of the other folks who are very active in the opensource community, looking at ways to go and get DNSSEC more widely used within other places.

But if there are others out there who have suggestions, please pass them along.

STEVE CROCKER: I see in the chatroom that Rod Rasmussen, who does speak for SSAC, says this is way outside of SSAC's remit [inaudible] may have worked in this space.

DAN YORK: There you go, Rod.

RUSS MUNDY: This is Russ Mundy and I wanted to just thank the asker of the question. I didn't see it come in at the time but it is an idea that has been pushed on for a while and Dan noted that the DNSSEC deployment effort that Steve and I pursued for a number of years did a number of things in this space and the type of thing that you're describing was in fact the direction that we were pushing. We had some successes and there have been some products that have come out that were explicitly pointed at ease of use, but the more of those that we get whether they're opensource or whether they're actual commercial products, the better off we are.

One of the big successes in the commercial product space was when Microsoft supported the relatively straightforward deployment of DNSSEC into their enterprise centric architecture.

So, opensource and commercial products are doing it, but more are always good to have and easy is I think the critical factor there.

DAN YORK: Yeah. Definitely. I also put in the chat for the folks some of the tools that are out there are available at the dnssec-tools.org site that, Russ, you're involved and Wes Hardaker is involved with and others are involved with as well.

WES HARDAKER: One little extra comment which is that security technologies in general have always gone through this transition stage of you need a lot of expertise to understand it and deploy it toward it's got to be automatic where users actually don't need to understand DNSSEC. So the real end goal is they should just have it, right?

A lot of people today don't even know the difference between HTTP and HTTPS. There was a time period where everybody had to remember to add the "S" and nowadays we've actually moved to the point with pinning and all sorts of other stuff where it actually becomes sort of automatic. That's actually our end goal, not to educate the end user but to take away the need for education.

DAN YORK: Well said.

UNIDENTIFIED MALE: I have one comment. Just to say that I keep my foot in other realms like Internet of Things and blockchain projects and they look to DNSSEC to help establish a baseline to build upon in these technologies that's trustable. We don't have …

When people switched from terrestrial to antennae to cable, they'd get a clearer picture and maybe some extra channels. So there was a consumer attraction and we just don't have that with DNSSEC that's going to drive people to get something that they don't normally get.

We also fail I think in the presentation of being able to show somebody that DNSSEC is actively there or not so they don't know if they're actually getting it or not.

But I think that these other technologies, like IoT and block chain projects, are going to help create an attraction for people to make sure it's there. Not just the security, the stability, resilience and the other things we typically focus on.

So I think it's a very helpful thing to look at … Consumer demand will drive that a lot. I see that at my registrar.

UNIDENTIFIED MALE: We're also seeing that … Victor [inaudible] I see in the chat. We're seeing a lot of that interest in the email side, too, which is great to see, from the use of DANE for securing email server connections, server-to-server connections. Sorry, Jacques, I know you were about to say something.

JACQUES LATOUR: Yeah. I said that we're out of time and we cannot compete with the next plenary session. We're being told that this is it.

I think this was a good workshop. I think we need to allocate more time for DNSSEC for the next policy forum, I guess. That would be the lessons learned. I'd like to thank all the participants, presenters who took time in this workshop and that's it.

WES HARDAKER: Thanks to everybody who made this happen.

UNIDENTIFIED MALE: Thank you.

KATHY SCHNITT: Thank you, everyone. You may stop the recording now. Bye, everyone.

RUSS MUNDY: Thanks, all. Bye.

UNIDENTIFIED MALE: Bye.

[END OF TRANSCRIPTION]