ICANN67 | Virtual Community Forum – DNSSEC Workshop
Wednesday, March 11, 2020 – 09:00 to 10:30 CUN

KATHY SCHNITT: This meeting is being recorded. For the panelists, please remember to state your name before speaking and please keep your phones and microphones on mute when not speaking to avoid any background noise.

As we are utilizing the webinar rooms for the virtual session, attendants will only be allowed to ask their question or comments through the Q&A box. Once you do so, at the end of each panel session I would be happy to read that question or comment aloud. This session will run from 14:00 UTC to 17:15 UTC with a 15-minute break from 15:15-15:30 UTC. I will now turn it over to Dan York with SSAC to begin. Please, go ahead.

DAN YORK: Thank you, Kathy, and welcome to the room, Kathy. I need you to drop the sharing so I can put the sharing on here. Okay. So, welcome, everyone, to our DNSSEC and security workshop. I'm Dan York from the Internet Society, based in Burlington, Vermont, an area in the United States. Welcome to you coming at us from all around the world as we do this virtual meeting in what we have.

Today, this session is brought to you by the Program Committee, many of whom are on this call with you, who are a group from across the world who gather weekly to talk about how we can do this to provide a program at these ICANN meetings that talks about DNS security, DNS privacy, and we've expanded this over the past year to also include routing security, which is why this is called the DNSSEC and Security Workshop. We'll have some pieces of those both, here.

This is an organized activity of the ICANN Security and Stability Advisory Committee, or SSAC. It is also with some additional assistance from the Internet Society, and the program are there. That is part of the [inaudible] of which this operates.

The program that we have today is I'm obviously speaking first, here, just to talk a little bit about what this session is and some of the counts we have and the things that we're doing. Then, we'll be moving into a moderated panel when Jacques Latour will be monitoring it and talking about the KSK Rollover and what it would look like in the future – all of those different kinds of things.

We will be pausing for a coffee break. That is part of how this is all working as ICANN has moved to this virtual environment. There are some breaks that are built into the session but we have arranged it so you will not have to disconnect. You will be able to continue to be on this Zoom session. You will not have to change off and go to somewhere else around that.

We'll then, after that, come back into this session where we'll have a 40-minute panel presentation about DNS over HTTP, otherwise known as DoH. As you can see here, we'll have a couple of presentations that we'll be talking about: "What is DNS over HTTP?" and some of the challenges/the pieces around that. Suzanne Woolf will be presenting about the upcoming SSAC advisory around DNS over HTTPS and DNS over TLS, DoH and DoT. So, there will be some of both of those being discussed there.

And then, we will conclude with a presentation driven by Steve Crocker where he's going to talk about some of the challenges that we've seen with DNSSEC provisioning with third-party DNS providers. This comes to some of the … Basically, if somebody is operating your DNS on your behalf, how do they update the parent registry, etc., with the new keys? DS records and things like that, that we'll get into as we go through this.

Throughout this time, as this program goes … Sorry, just to end. We'll end that with a little bit of a review at the conclusion of this, just to also get some of your feedback about how this worked. This is obviously our first time doing a fully virtual environment, although we've always had remote participation.

As Kathy is putting in the chat – which if you are not able to find it or if it's not there, if you go up to the top of your window there will be a little box that drops down and there's a chat thing that you can see there. You can also go to the three dots where it says "more" and you can say "exit full-screen mode" or somewhere in there and be able to get out of that if you'd like to have the chat box next to this. In the chat, you can go and do that.

You are able to go and submit a comment there. Kathy and others will be in that space. You can ask a question or you can leave a comment in there, as well, if you'd like something to be read out loud, in there. Those of us who are more purist may be wondering why there isn't a slash before the N-tag, but whatever. We'll survive.

Okay. I'm going to begin by just talking a bit about the current status of DNSSEC deployments around the world and some of the statistics and things that we've seen. To begin with, we look at the DNSSEC side. We look at the two sides because there are two parts to DNSSEC. There are the signatures, the signing of domain names that is done at the authoritative server level, and then there is the validation, which is the checking of signatures. That's done at the recursive resolver, as what we think of as the DNS servers that we use all across the Internet for all of our normal browsers, and applications, and everything else, mail servers, anything that's doing it.

But what you're seeing on the screen is the DNSSEC validation graphs that APNIC labs generate. Geoff Huston and George Michaelson have been on our sessions in the past to talk a bit about how they do this. They have a whole testing infrastructure that goes out across the world and tests whether or not people are able to perform DNSSEC validation from their networks.

What's great to see here is that we're seeing, now, a fairly steady growth happening over this past while to where almost a quarter of all DNS queries or all DNS locations for all networks—I guess we could say "across the Internet"—are being able to do DNSSEC validation.

Now, obviously, it's significantly higher in some parts of the world and it's significantly lower in others, but this is great to see, this kind of continued growth happening as we've gone through this space around this.

We can see here, also, from the APNIC stats, the list of which regions of the world have the highest percentage of DNSSEC validation happening. You can see some of it in Asia, Oceana, and some of those areas are having the largest percentage, according to the APNIC stats, of where this is actually happening.

Now, on the findings side, we've also seen some fairly significant growth over time in the number of signed records. These are statistics coming out of the DNSSEC-Tools project, where Wes Hardaker, Viktor Dukhovni, and some others have been working on a number of statistics packages to go and see what's happening here.

You can see, again, a nice trend in overall growth of DS records. And so, these are signed domains that are there. So, this is the number that you're seeing continuing to grow up, coming up nearing 11 million. I mean, obviously, it's still a small percent of the overall number of domain spaces but it's a great kind of growth that we're seeing coming out of this.

Also, one of the things that we have been promoting here is the use of DANE records, which on a simplified level are a way to put a TLS certificate or a fingerprint of a TLS certificate into DNS, sign it, and be able to assert that that's the certificate that you want to be using for a given domain, for a given type.

Viktor Dukhovni has been doing a huge amount of work over time tracking the number of DANE records being used for mail exchange, for MX records, for sending secure e-mail. You can see here, starting in early 2019, we started to see a large number of deployments and that has continued on, now, into 2020, where we have a very strong ongoing deployment of records that are DANE records, which will enable people to send signed and secure e-mail.

On the RPKI side, which if you're not familiar is the Routing Public Key Infrastructure. It is a mechanism to go and assert that you have to right to originate this route, essentially, is a way to talk about [that are here.] This shows us the growth in the number of people using RPKI and doing validation. In this case, again, it's the validation side of it.

You're seeing, again, a growth from 17% in November, now up to 18-and-a-bit percent/almost 19% here. Again, we're seeing growth in this regard. We're also seeing a nice growth in the number of people who are signing their routes and putting those out there in different ways. This shows a growth across the different RARs and what's happening out there.

Again, nice. This is exactly what we want to see, growth in both the validation and in the signing for both DNSSEC and for RPKI. So, good records to see around that.

I just want to put out a final point. We have been maintaining for a good number of years a set of maps that show the deployment of DNSSEC around the world with ccTLDs. We're continuing to see a good number happening across the space. And so, one of the areas where we do need to see more growth, at this point, the CCDS that are remaining to be signed are primarily in Africa, a few parts of Asia, a few parts of Latin America. Overall, we're continuing to see good growth in the signing of ccTLDs.

With that, I want to just wind down with a couple of different resources, to say that there are a number of them. We at the Internet Society provide some resources at the Deploy360 program, which is a program we've operated for a number of years that has a number of different things. I put a new URL on here this year in 2020. I am leading a project at the Internet Society called Open Standards Everywhere, which is designed to focus on how we promote the use and deployment of open standards, with a specific focus on 2020 on web servers.

We're looking at, "How do we raise the level of security and availability of web servers?" and specifically around looking at how we get more web servers using IPv6, DNSSEC, and TLS. At that internetsociety.org/cse it jumps you to a page where it has some information about this project. DNSSEC will be part of that work and we will have some information in there pointing to many of the other resources that are out there, as well. But it's something to just watch over this over the year. If you're interested in learning more, you can find out at that URL.

Also, DNSSEC-Tools has a good number of resources that are there. As you say in the statistics, we have some nice stats that are part of there. We also have the stats from APNIC. Also, for more historical information you can go to dnssec-deployment.org and you'll see information about the history of things.

There are also RPKI resources. I should be clear, too. I misspoke earlier. It's not the Routing Public Key Infrastructure. I apologize for that. It's the Resource Public Key Infrastructure. Anyway, that's what RPKI stands for. It is there and there are some good resources, as you can see here from NIST, there. Also, some stats from RIPE and securerouting.net also provide some other information, as well.

With that, I am going to turn it open to questions. Before we get going I will just stop my sharing and ask if there are any questions that people have before we begin the show. Any questions about how we're operating? Anything that we're doing? Now is your moment. Put it in chat. Let us know. I'm pleased to see right now that we have about 86 people online, which is great to see.

UNIDENTIFIED MALE: And no comment so far.

DAN YORK: And no comments. I guess everybody must have understood everything I said or something like that. Well, then, with that I'm going to thank you very much. I look forward to your participation in the rest of today’s sessions. Please, again, know too that at the end we are going to open this up for questions. We’d love to hear your comments. How did this work? What comments/feedback would you have about how we could do this differently, better, whatever else, and pieces like that?

Oh, we did get a question from Paul. The question is, “What accounts for the rise of gain in RPKI deployments from 2019 onwards?” Well, I’ll say mine, and panelists, I’ll be able to open this up to you all in a moment. I would say this: one of the things that, particularly on the DANE side, was some very focused efforts on deployment by some of the large e-mail providers and some incentives/some programs that people were running, particularly in Germany and particularly in some others.

There was some guidance issued by the German Security Information Ministry encouraging everybody to use DANE for securing e-mail. I seem to recall the Dutch National Security did that, as well. NIST came out with some guidance around that. So, there were a number of different information security agencies that strongly encouraged people to do that, as well as within the industry itself a number of the providers working with that on DANE. With RPKI I think, again, it was also some deployment efforts within the RARs. Russ, do you want to comment on that, perhaps?

RUSS MUNDY: Thanks, Dan. Yeah. I think in addition to those, which were all very helpful, the fact that some of the open-source mail packages have a capability embedded in it, now. So, there is a lot of similarity between what happened in the early days with signing of zones once the regular open-source vendors started to include it and include tools to make it easy to use. People, in fact, started to make much broader use of it. So, I think that's the other factor besides the policy and promotion activity.

DAN YORK: Yeah. I think another thing, Paul, too, that we saw was there were some efforts within the larger industry. For instance, the manners, the MANRS, the Mutually Agreed Norms for Routing Security effort, has been going on for a while to facilitate an expansion of routing security. That effort has grown dramatically over the past couple of years.

There are now up to almost 300 network operators and I don't even know how many actual networks/ASNs are involved, but a substantial number of people who are agreeing to abide by those norms, which includes the use of RPKI for routing security. So, I think some of those concentrated deployment efforts really helped move that on. Any other panelists want to speak to this point?

WES HARDAKER: Yeah, I can speak into it a little bit. As you know, Viktor Dukhovni is the one who has been collecting good stats and things like that about [end] deployment. He has actually had a lot of conversations with both software implementations as well as major providers that do SMTP service on behalf of domain owners. And so, there has been a couple of large jumps in DANE uses just as those major providers suddenly enable DANE for SMTP as mail support across all of the domains that they have.

They received sudden jumps in both DNSSEC signing as well as sudden deployment in DANE records. So, if you look on stats.dnssec-tools.org, you'll see huge jumps in the very top graph. One of the biggest ones was one.com, which was a major provider of SMTP service, basically, signed everything all at once. And so, there was this huge rise and it's all from one vendor. He has been the promotional aspect of a lot of it. We owe him a lot of thanks for that deployment.

DAN YORK: Great. And we do see a couple more questions, here. One, "Please let us know how the hosting providers can enable DNSSEC from their C-panel, WHM, or Plesk." I don't know [cross talk].

WES HARDAKER: Yeah, but he's asking if he has a domain in a C-panel on some site. "Will that hosting provider provide a check button or not?" Jacques, I agree with you. It should be default but it's not. He's wondering how he can turn it on, and unfortunately, I think you need to probably contact your support contact for that provider that is actually hosting C-panel for you.

DAN YORK: Yep. I see another question here from Imran about, "We are a web hosting provider. We configured our webservers to IPv6. Now, we're facing some issues with SSL and Comodo. Should we deploy more servers with IPv6 with these common issues, or what is your suggestion to solve this issue publicly or globally?" So, I'd be glad to correspond with you directly on that one because I can connect you with some people who can help with that.

With that, I think if we have no further questions I'm going to turn it over to Jacques to begin our next session around the KSK Rollover. Thank you, everyone. Please continue to engage with us throughout this day. Thank you, all.

JAQUES LATOUR: All right. Thank you, Dan. Good job. It's a virtual clap. Next up is our panel discussion on the KSK Rollover. We have about 45 minutes on this. First up is Wes. He's going to talk about the roll, roll, roll the root. Go for it.

WES HARDAKER: All right, thanks. Let me share my presentation and you guys have to tell me if you can see it. I think that window. Does that look about right?

JACQUES LATOUR: Yep. Looks good.

WES HARDAKER: Oh, they've put me out of full-screen mode. Let's see if sharing it in presentation mode actually works. It may or it may not. I don't think it does. All right, that's fine. We'll go this way. So, this work has been presented in a number of other places, so for those of you who have already seen it, I apologize.

However, it has not been presented at the DNSSEC workshop. We were going to present it six months or so ago but this was a paper published through the IMC, Internet Measurement Conference, and they decided to go completely anonymous review, meaning you couldn't really present it anyway else or else it might have affected the reviewer's review of the paper who didn't know who wrote it.

It was a collaborative effort by a large number of people, as you can see on the bottom list. We actually had seven people that worked on this. Really, we wanted to tell the story of, "What was the KSK rolling process like for the DNS?" So, we will dive into the details.

I think, as most people here know because this workshop has been running for such a long time, DNSSEC brings integrity to the DNS. In particular, validators need to be able to have a public key in order to resolve stuff, so they have to have a trust anchor. So, as we look through points, the root zone has a key and the key data is used to sign the data that is returned to recursive resolvers.

And so, the [ADA] gTLD servers, for example the servers for .com, they sign those and the resolver has to have a matching key. If the resolver and the root both are using the same key then the resolver can verify the signatures given by the root.

So, what happens when they disagree, when the keys don’t match? That’s where we run into a problem. So, what happens when you roll the top key? If you roll the key at the root then all the resolvers have to update. This was, of course, the conundrum in 2017 and beyond.

And so, what happens is, when there is no key that matches there is no validation. End users get no DNS responses. Basically, resolvers full stop returning any responses. And if 30% of Internet users rely on validating resolvers … As I think Dan just showed, it’s actually more like 25%. Various surveys have different levels of measurement. Every validator needs to have the current key signing, which is the 2017 version of the KSK.

Validators use hard-coded keys a lot of the time. Sometimes, they’re actually written into the code so if you haven't updated your software you might have the wrong key. Dockers containers and similar have a challenge because they don’t have writable storage so the key updates may not happen.

In general, DNS, since its inception, has been pretty much a fire-and-forget type of technology where once you deploy a resolver you really don't need to touch it again. Even authoritative servers, to a large extent, you don't have to worry about touching data again.

DNSSEC very much changed that. It requires resigning on the publication side and requires the validators to get updated keys and to watch error messages at a higher level than before.

So, this is the timeline for the key rollover, the KSK 2017, which is what we will call "the newer key." The 2010 was the older key. It was first published in July of 2017 with the intent that it would roll over sometime in October. That was halted, and we'll talk about that stop in a minute. It was halted for an entire year, and then it was resumed later in September of 2018.

And then, the rollover actually happened where the material published by the root zone was actually signed with a new key, starting in 2018. And then, the older key, KSK 2010, was revoked on January 11th of 2019, last year. And then, it was finally removed entirely from the root zone in March of 2019.

So, I'm first going to concentrate on what happened before the rollover, so what happened leading up to the events of the actual day of the rollover that happened in 2018.

A couple of notes. There are a few ways of trying to measure this information. RFC 8145, which has been discussed in the DNSSEC workshop in the past, was a mechanism where resolvers were supposed to signal to the root zone servers that they were occurring what keys they actually had in storage/what keys they had available to do validation with.

It was sort of the goal of guesstimating the number of resolver validators that had the new key. It wasn't as entirely fantastic as we had hoped but it was the only signal we had at the time. We used data in the paper from A, B, and J roots. So, that's Verisign's roots and USC-ISI's root. We looked at signals for up to 100,000 validators that were sending these 8145 signals daily.

So, during the uptake of KSK 2017 after it was added, which was in July of 2017, where the dotted line indicates, KSK 2010, the red line is the number of resolvers that had the old key in their storage. So, we're looking at the number of resolvers that were transmitting the 2010 key, what percentage of them were also transmitting the 2017 key saying, "Yes, we have this one, now."

You can see that 30 days or so after the key was added there was a whole down-period where they weren't supposed to trust anything new for a while. All of a sudden, there was a sharp uptake where the number of resolvers that had the new key was large.

Note, however, that there is a gap at the top and that gap caused a whole lot of concern because 8% of the resolvers didn't have the new key. Well, 8% of resolvers is a fair amount, so ICANN kind of threw out a, "Okay, let's stop. Let's hold. Let's figure out what's going on," because they were really hoping that the number of resolvers would be closer to 100% with the new key.

So, a number of things have been where some research was done in order to zoom in on, "Let's see if we can figure out what that 8% consisted of, where they had problems, and see if we can figure them out." So, one of the things that we did was we looked at the sources of the signals and looked at a number of the resolvers that only sent a very small number of queries in a very short period of time and then suddenly went silent. They'd come up, they'd advertise they only had the old key, and then went away.

And so, the query for the old key that says "I have the old key but not the new one" is that top one out of a bunch of DNS packets that were received. There are 15,000 saying "I only have the new key." Of those, the next most popular query that they sent was, for a period, the root. They wanted information about the root.

And then, the next two most popular queries that were sent came from VPN domain providers. In other words, there was a sudden belief that, "Oh, well, if the most popular domain that these particular addresses were sending were for our VPN provider," that might be the source of the problem.

Sure enough, that was the source of the problem. So, here's a graph that was later on from the previous graph where these are the number of signalers that only saw the 8145, that only saw the old key, and their IPv4 and IPv6 signals. As you can see, that's actually closer to 20%. So, by the time that these measurements were actually started, that 8% that occurred before was now up to 20%. So, the percentage actually got worse, not better, over time.

All of a sudden, there are these drops on the graph because we contacted the VPN provider and told them of this problem. Sure enough, they've verified that they had hard-coded only the old key and when they made the software releases we saw remarkably sharp drops in the percentage, which made us feel like we'd found a pretty big source of the problem. You can see that the graph is rising up on the left and it's slowly, steadily, going down with fairly big jumps downward on the right. So, that was a decent finding.

After this and a number of other conclusions, ICANN decided that the rollover should happen and, sure enough, on October 11th of 2018, they did switch to using the new key entirely.

So, during the rollover we are going to switch to looking at, what was the user's perspective? What actually happened at that rollover point? And so, as most people know, RIPE Atlas is one of the easiest and most prevalent ways of trying to measure how users perceive the rollover.

And so, the approach was that we would use RIPE Atlas probes, send measurements once per hour, and then see if they had cached KSK 2017, and if they validate correctly. So, as you remember, RIPE Atlas probes hopefully sit behind and use real resolvers. They're not doing resolution themselves but they send their queries to resolvers so that we can be able to measure the resolvers that were actually deployed in the real world. With that, we measured 35,719 resolver addresses from 3,141 different ASes, networks, or ISPs.

So, when the KSK 2017 was activated the vertical black bar on the left was the timing of that. Remember that data in the review team has CTOs, generally, of two days long. So, it takes possibly two days before resolvers will fully take up the new data because they don't query that frequently. That the data is useable for two days so we expected a two-day window and that's exactly what this graph shows, where at about slightly less than 24 hours we actually see the 50% point flip and the number of resolvers using the new key and having it cached surpassed the number of resolvers with the old key having it cached.

There are a number of large, jagged jumps. You can see that the line is fairly jagged and there are jumps upward and downward where fairly large resolvers suddenly started validating with KSK 2017. That may have been seen from multiple points of presence. It was a very good indication that the validation and resolvers deployed in the world were actually having a success because you know that one goes to zero and one goes to 100%, and that's exactly what we wanted to see.

There were, of course, some exceptions. Of those 35,719 resolvers, 34,000 were always secure or always insecure. They were always returning the proper results according to their particular configuration, but some weren't. So, some 970 were secure beforehand. In other words, they were validating beforehand and then suddenly went bogus. 747 of them were secure beforehand and then suddenly went insecure, which could be because they actually turned DNSSEC off.

519 sent a whole bunch of excess DNS key queries, which we'll dive into in a little bit, 359 of which sent 1.5 more times DNSSEC key queries after the rollover, trying to find the right key. In other words, they were trying to look for the keys and look for signatures and they appeared not to have the right key in their configuration.

218 of those were fixed within an hour. In other words, there were probably 218 different ASes that went, "Oh. Uh-oh. Something's wrong. Our entire DNS infrastructure is not working," and they realized what the problem was and fixed it somehow or another.

138 were fixed after an hour and then three were never fixed, and it's possible they had been fixed by now. It's entirely possible that those three just weren't used, they were secondary DNS servers or something like that. They just never noticed. It's possible they've been fixed by now. This paper was published again last year.

EIR was one of the biggest ones. There was eir.ie, which is in Ireland. It was the major ISP that seemed to be down for a fairly long period of time. I don't remember the actual timing but they were down, I think, for hours or something like that. They said it was related to DNS. I don't believe that we've ever had contact that has confirmed it was related to the key roll but it's certainly likely.

Taking about the outages in general, we did see a number of strange things happen. Right after the rollover, there was a gigantic spike in the number of queries to the root zone for DNSSEC-specific information right after the rollover. That's likely caused by validators not having the right configuration and suddenly sending a lot more queries toward the root for DNSSEC-related stuff. EIR would certainly have contributed toward that spike if that indeed was their problem. That spike is not all EIR, however. At least, we don't think so.

There are a few other strange things that we're going to talk about. If you look at the rollover, there was a large spike there. That was sort of to be expected because we knew some people would have problems. After their revocation, so the revocation that was set in the key and republished in January, and then later, when it was finally removed from the root zone, there was a mysterious bump after that. That's actually still sort of unknown as to why that occurred at this point.

But we're going to talk about that end-time period before, so we're going to talk about both that revocation point that I just highlighted, as well as when it was removed from the root zone and what traffic looked like before and after that.

So, if we go back to that same sort of graph, there was an increase after the rollover of DNS key queries. So, this is people querying for the root zone, saying, "What keys do you have? What keys are in your zone?" So, point one, there was sort of an expectation that there was a partial increase expected right after the rollover. Nobody was shocked that it went up a little bit. It would be nice if it didn't go up at all but nobody was that surprised that there at least there was a little bit of an increase. Some of that could have just been from monitoring for that point.

Revocation. After the revocation point, however, there was a very sudden expected increase. So, when the old key was transformed in the root zone with the revocation bit set, all of a sudden there was a major increase in the number of DNS key queries being received at the root zone, to the point where at the peak, just before the removal, 7% of the total query load—and I think that was from A and J-root—was from DNS key queries from, clearly, resolvers with some sort of problem.

There was a big question of, “Well, what happens at the removal? Does it get better or does it get worse?” That was a good question. We were lucky that it sort of returned back to normal right after the rollover, which is a very good thing.

So, what was behind those query [flex]? What causes all of that to happen? One of the things that we did is we sent chaos queries, which is the mechanism for determining what resolver software was on those sources that were sending too much data. That determined that it seemed to be some older versions of BIND that were causing the problem, that had a bug in it.

And then, we also did some outreach. OBH was one of the sources and confirmed that they were running by 982 on CentOS, which exuded that big that was causing queries after a revoke key.

And then, we reached out to Perdue University and confirmed that they had a DNS lab and actually were able to provide us with the BIND configuration that was causing this particular incident.

So, we took that BIND configuration from Perdue and then tried to reproduce it in the lab. It was actually kind of hard to reproduce because you ran it and it looked just fine. So, we ended up running it 20 times in a row with 20 different stop, BIND, replace the configuration, restart it, and see what happened.

In a number of conditions, when the DNS keys managed/contained 2010 but not 2017, the DNSSEC enable side was set to “false” and the DNSSEC [inaudible] side was unset, leaving it in a state of “yes.” When you run that 20 times, the number of queries sent during that experiment, you can see that there were spikes on experiments 7, 13, and 17 where they suddenly just started ramping up and sending a whole bunch of queries, meaning after 20 different experiments it wasn’t consistent from time to time. So, something internally to the BIND software was just randomly triggering into a weird state where it was sending way too many queries and those queries burst only occasionally.

After the revocation, when we also studied the result of, “What keys did validators actually have available to them?” after revocation you would expect, since that revocation was published, that bit was published in the key saying, “You should no longer trust this.” People stop trusting it.

And so, you can see that the red line dropped down to … I'm not quite sure where that line is. It looks probably about 12%, where probably some validators had it configured to trust it so they did. The other ones that were managing it actually dropped it out of their trust zone. All of a sudden, it started creeping back upward. You can see that over time, since the key was removed, some root validators started thinking that it was an okay key again, even though ICANN should never be using it to sign the root zone again. It has crept back up ever since then.

Partially we have determined that it was caused by some older version of Unbound. We believe that a lot of people are distributing configuration. As they update software, they're redeploying their existing configuration with the older key. That caused validators to believe that the older key was still possibly going to be used.

So, that's the end. There is about ten minutes left for discussion and questions. I do have some back-up slides if we're bored but that's essentially how it went. Wait a minute. Where there was the conclusion slide? Oh, I'm sorry. There are a few more slides.

Do we need to improve telemetry? So, a couple of takeaway points. 81, 45, and 85 over nine are the two telemetry things that give us the feel for how well deployment, and validation, and the KSKs are existing out in the real world. Specifically, we need to be able to identify the true source of the signal.

One of the problems is that we don't have an idea, once we get a signal, of, "I'm a validator and I have this key available to me." We have no idea how many users that validator is actually serving so it's very hard to get a feel for real-world deployment. It could be one validator with one person behind it, which is actually the case with the VPN software, or it could be one validator with millions of people behind it that are serving an entire country.

There are privacy concerns with telemetry. Any time you're transmitting information you're sort of giving away information about yourself, which is both good and bad.

So we need to change trust anchor management? I think one of the essential takeaways from this rollover experience is that when you shipped trust anchors generally in the operating system there are far less places to update. There are far less operating systems out there than there are trying to get everybody to install keys automatically or using the 50/11 process which is dynamic management. It's much easier if we can ship keys ahead of time in software, especially in operating systems.

Finally, first off, even though this work was presenting some of the issues that were seen and some of the ordinance to it, I think that the rollover was generally considered a success by most people.

Independent analysis and measurements are sort of necessary on the Internet. You know, the who Internet Measures Conference that this was submitted to is designed to do just that: how are we going to measure things in the Internet?

Though I gave this big, long stream of measurement as one thread it's important to point out this took a lot of people. There are seven authors to this paper with a bunch of institutions that put this work together. It was a collaborative effort to bring this complete story to light.

And then, we really need to keep telemetry in mind. I think in any new protocols being developed, DNS or otherwise, we have to have good ways of doing this measurement. As I've already said, trust anchors should really be managed centrally when possible. Questions, suggestions, comments?

JACQUES LATOUR: Thank you, Wes. Do we have any questions for Wes?

KATHY SCHNITT: We have a question in the Q&A pod from an anonymous attendee. "At [inaudible] and certain domain registers. Namecheap have an option under 'guest' where it allows you to activate DNSSEC for the domains you own." That was a comment.

WES HARDAKER: Yeah. There are a number of providers that are doing that and one of the biggest … In fact, I think GoDaddy is still the biggest registrar out there in terms of the number of domains registered. They will be enabling a button to do that in the future. Fantastic new, because I think they corner about 40% of the domain market.

And so, when they do that, I encourage anybody that has domains registered under them, or if you have Namecheap or other stuff, go click that button. There should be no downside to do that.

KATHY SCHNITT: At this time there are no further questions or comments.

WES HARDAKER: Okay, then I will turn it back to Jacques.

JACQUES LATOUR: Thank you. The virtual clap.

KATHY SCHNITT: Yay.

JACQUES LATOUR: All right. Next up is Kim Davies from IANA PPI/ICANN, who is going to talk about the KSK Rollover update plans.

KIM DAVIES: Thanks, Jacques. Can you hear me okay?

JACQUES LATOUR: Yes.

KIM DAVIES: Perfect. So, I have a slide to present to you an update on the consultation we presented through the last meeting regarding how we will do future KSK Rollovers. Next slide, please. So, I'll give you a brief update about that but I also wanted to seize the opportunity to give you an update on some operational matters that have happened just in the past few weeks. This is the right audience to share it with so apologies for usurping the slots slightly for a different topic. Next slide, please.

A little bit of background, although I think some of this duplicates what you just heard in the last presentation, the first KSK was created in 2010. Part of the design of KSK in general was to make use of design teams that were constituted from community experts, folks from our operational teams, folks from Verisign and so forth, to drill into the different details and come up with a set of recommendations.

We did this prior to creating the first KSK. There was a design team that built the original design of the root zone DNSSEC setup. We also convened a design team with respect to how we should perform that very first rollover. There was a design team report that developed a set of recommendations and those recommendations were used to implement the first rollover.

So, the first rollover was originally scheduled for 2017. As you heard, there was a pause to consider some of the telemetry data that we were seeing, but ultimately that KSK 2017 was put into service in October 2018.

You heard in the last presentation, I think from our perspective it's the same. We consider that the rollover was performed successfully and that there was minimal disruption as a consequence of the rollover.

So, we had a design team that designed a process. We implemented that process. As a consequence, we had a relatively successful rollover. [inaudible] we want to do now. Next slide.

Internally, we were looking at a lot of the discussion that was happening in the community and wanted to capitalize on that level of interest that we were seeing so that a lot of folks' thoughts and impressions on the rollover process were captured.

So, the first thing we did was to solicit comments immediately and direct them to be captured onto the KSK Rollover list for capture. I'm seeing comments that I'm breaking up so maybe I'll just pause my video in case it's a bandwidth issue. Okay. Carrying on.

So, we directed those comments to the KSK Rollover list initially. We wanted to make sure those that wanted to contribute their feedback could do so in the moment. And then, we also undertook to analyze those comments in the second half of last year and produced a recommendation for future rollovers.

Some of the common themes that we got from that early feedback that we captured was that rollovers should be routine. There were actually a lot of comments that said it should be an annual event. There were comments suggesting we explore back-up of standby keys, perform more monitoring of the impacts of large key sets, and consider alternate signing algorithms such as ECDSA. These were all suggestions in that early feedback that we took into consideration. Next slide, please.

So, our proposal that we put together tries to create a predictable approach to future rollovers. We discussed the benefits and complexities of different rollover intervals and we proposed a three-year rollover interval. This really was designed to balance the desire for a relatively routine and regular rollover on a predictable schedule but also recognizing there is a lot of hidden complexity in doing the rollover itself.

Certainly, we felt that an annual rollover process was far too complicated for us to enact based on the way we do ceremonies and key management today. If we wanted to preserve the fundamental design of KSK management it didn't seem that annual rollovers as a routine event seemed like the optimal balance to strike.

Our proposal does extend the amount of pre-publication for new keys and new trust anchors. The idea there is that we would provide a greater opportunity for propagation before the rollover. You heard in the last presentation the desire to provide a capability for things like software vendors/operating system vendors to bake-in the trust anchor set in their software distribution practices, and this would provide a greater window to do things like that.

But one thing that we've retained is today we use a phasing approach that revolves around calendar quarters and key ceremonies. We've tried to align with that, as well. Next slide.

Okay. So, we published an outline of the approach and put it for public comment—as many of you will recall, we presented it at this very meeting at the last ICANN meeting in Montréal—and put it for public comment. Based on the feedback we received, we extended the public comment period to the end of January. We received 11 comments, that public comment period.

Current status on that public comment period is that we are now due to compile a staff report that distills those comments. For reasons I'll get to in a moment that has been a little bit delayed but we have done a cursory review of some of those comments. Nothing too surprising there, and I just wanted to briefly go through some of the common themes we saw in a number of comments. Next slide.

With respect to pre-publication of the standby key, we did see various opinions on whether it was problematic. Again, we'll go into a deeper analysis of those comments and produce a staff report soon.

There were various different suggestions relating to changing the exact timing of events within the process. One such suggestion was tweaking or modifying the number of keys or the phases so that there is always coverage throughout the calendar with a standby key.

The proposal we put out for public comment provided that coverage with only around two-thirds of the time. So, the suggestion was, "Can it be modified to provide constant coverage?" There were a number of suggestions that we consider explicitly Sunset provisions based on the key strings. There was one comment, "We should be mindful of a skill loss if the events are too spread apart," and various other suggestions for tweaking the phasing.

There were some risk mitigation-related comments, such as, "If we generate the standby keys, they don't necessarily have to be kept in the HSMs. They could be exported and reimported at a later date," and another suggestion relating to a different potential mechanism for generating the key itself.

I would say that, probably, one of the most common threads in the discussion was that of an algorithm role. Our proposal was, essentially, to separate consideration of the algorithm role into its own unique consideration as opposed to integrating it, necessarily, into the process of regularizing KSK Rollovers.

Some responses suggested that some felt it was necessary that this be integrated so it would kind of be a blocker for progress on this project. Others indicated they felt it was okay that it be a parallel activity. Next slide.

There seemed to be general support for the cadence, i.e. doing it every three years. There were suggestions relating to how we perform outreach to communicate activity in support of the KSK Rollover.

There were quite a few suggestions that I would classify as editorial, different ways we could convey, essentially, the same information.

And then, there were some ideas for where IANA might play a larger role than it does today or historically. Some of the suggestions were IANA using more measurement. It was unclear whether that is performed directly by IANA or relying upon the measurements we've heard about already during this session that are performed by third parties. Also, things like whether IANA should be providing advice to regulators and policymakers. Next slide.

One thing that was a bit of a surprise to me, at least, was that the common theme in a number of the responses was an argument for a lot more specificity in the consultation material. Our thinking in presenting this for public comment was to get high-level agreement to the principles that we had outlined and then, assuming that there was general community consensus that this was the right approach, we would then go into, obviously, the detailed business of producing an implementation; what the exact DNSSEC practice statements would be, what the new operational practices and procedures would be. So, to take the guidance and turn that into an implementation.

Really, our objective here was to capture what worked well and capitalize on that from the first KSK Rollover. We weren't necessarily trying to boil the ocean and try and resolve all of the open questions around KSK operations.

I think, in reality, a lot of the more challenging issues that are still unresolved for KSK operations, particularly thinking of the algorithm role, really do require a lot of significant research that is simply beyond IANA's competency to do directly. We would need to recruit resources from elsewhere in the ICANN Organization, in the community, and so forth, particularly where it involves funding to support that. That needs to be planned with significant lead time.

So, this is significant work from my perspective and we wanted to make sure that we could make forward progress on the KSK Rollovers without necessarily addressing all of those issues.

It's not clear if this approach, based on the comments that were received, was communicated well enough, and it's not clear, to me at least, whether it is just a matter that that proposition wasn't communicated well or if there is a fundamental disagreement that that is the right approach. Maybe, there is a feeling that a lot of these details aren't necessarily a prerequisite to get community sign-off before we go into the more detailed implementation work. Next slide.

So, that's all I had on the KSK Rollover activity that we're doing but I did want to segue into some operational updates on the KSK ceremonies that we perform.

The first one I wanted to update you on is one that we recently concluded a few weeks ago, KSK Ceremony number 40. Next slide. So, KSK Ceremony number 40 was originally scheduled for February 12th, 2020. The objectives here were, as is typical of a key ceremony, to sign one calendar quarter's worth of key material. In this instance, it would be a key material covering April through June of 2020.

We also had a maintenance objective to decommission an HSM that we had that is no longer in use but we also had some pre-ceremony activity to perform some maintenance on the facility. This maintenance was specifically in relation to upgrading some of the lock assemblies within the safes that we have in the key management facility.

Our philosophy in this maintenance work is to perform what we call "Administrative Ceremonies." These Administrative Ceremonies are audited to the same standard as the key signing ceremonies but they do not involve actually activating the sensitive materials within the safe, i.e. using the HSM, using the smart cards, and so forth.

What we use Administrative Ceremonies for is, obviously, maintaining the physical equipment but also inducting new staff members into the trusted roles, for example.

We hold these Administrative Ceremonies adjacent to the public Key Ceremonies and we do invite the trusted community representatives who are available to witness these Administrative Ceremonies.

So, apart from the fact that they do not actually perform any key signing activity, these ceremonies are thoroughly witnessed and/or audited, as well. Next slide.

So, on the 11th of February, as I mentioned, we were doing this pre-ceremony work to upgrade the lock assembly. In brief, we were not able to open the safe. The safe was being opened by a trusted role called a "safe security controller." They have knowledge of the combination which no one else does. They were able to dial in the combination correctly and there is a digital display on the safe that indicates that it was performed correctly, and that was indicating, but the actual, physical bolt within the safe door did not retract to allow the safe to be opened. So, this, essentially, meant that there was some form of electrical or mechanical failure within the lock assembly itself.

Irrespective of the reason for the lock failure the remedy is the same, which is that we needed to drill the safe. This is one of those scenarios that had always been contemplated as sort of a worst-case disaster recovery scenario. We found ourselves in a position of actually needing to exercise it in order to replace the lock assembly, to gain access to the safe, and to replace the lock assembly.

This work took around 20 hours to complete. Those 20 hours are spread across two days. This work involved, predominantly, drilling into the safe body and into the lock assembly, removing the bolt that wouldn't actuate, and allowing the safe to open, then replacing the lock assembly as well as remediating the safe back to its original condition.

Part of this work was complicated that the lick itself has a number of anti-defeat mechanisms and they were triggered during the drilling process, partly due to the construction of the safe itself. Next slide.

So, the conclusion of this was the ceremony was successfully conducted, albeit with a four-day delay. I think that the valuable thing that came out of this was we gained a lot of experience with how to drill the safe that we didn't have prior. Thankfully, it had no operational impact that we needed to do this, and that knowledge that we gained will inform our future plans for disaster recovery in other scenarios.

I wanted to acknowledge in particular the community volunteers that were there as well as our team, our staff, who worked literally around the clock throughout the Tuesday through the Saturday that this work was ongoing. A lot of people change their flights, changed their plans, and so forth, to enable it to happen, so we were able to successfully complete all the work that week. There was definitely a risk that this work would need to continue into subsequent weeks due to not having sufficient personnel available, so having everyone move their schedules worked out quite well.

We are making revisions to the way we conduct Administrative Ceremonies moving forward. The idea there is giving greater transparency. Next slide.

So, I wanted to, finally, finish with an update on the next key ceremony that we have scheduled. Next slide. So, key ceremony number 41 is currently scheduled for the 23rd of April. This is, arguably, the ten-year anniversary of us doing these ceremonies and just shy of the ten-year anniversary of a signed root itself.

The objectives this time will be, again, to sign a quarter of key material covering July through September, and also replacing two of the trusted community representatives who are retiring from their roles and being replaced with two new volunteers.

Currently, we expect this to happen as scheduled and as planned. However, with the coronavirus situation, we have focused our efforts on developing contingency options in case the situation deteriorates.

Some of the work that is ongoing right now is periodically revaluating the situation of all the participants, ensuring their ongoing ability to travel, and continuous monitoring of the threat as it evolves, and building out those contingency scenarios.

One thing that is useful to know that was in the original design of the key management facility is that it should be possible for a staff-only ceremony to be conducted as a disaster-recovery scenario. So, whilst normal operations necessitate we have a certain quorum of community representatives attend from all around the world, in the event that international travel is substantially halted we do have options to perform key ceremonies absent that attendance but some of the exact triggering conditions around such an event have not been well defined. Next slide.

Some of the contingency ideas we're working up roughly in the order of severity from least severe to more severe. We can hold the ceremony with less than the ideal number of people present but still maintain our minimums. We try to have a lot of overlap in the roles in case someone has a flight delated, someone gets sick, someone forgets something, and so forth, but we can hold it with less personnel.

We could advance the ceremony data and have it sooner on the notion that the sooner we hold it the less likely the impacts on travel will be. We could postpone it in the event that the selected day is not viable anymore but there is a reasonable prospect that it could be held at a later date. We could hold the ceremony in the alternative facility. This ceremony is due to be held in the US East Coast but we have another facility in the US West Coast that we could use.

If there are specific trusted community representatives that are unable to travel due to their geographic location we could potentially induct new TCRs to replace those that are unable to travel.

One idea to help mitigate impacts down the road is potentially signing key material beyond a single quarter. I mentioned that we sign material for three months but it could be for longer.

We could perform ceremonies with less than three TCRs physically present, that’s the option I mentioned earlier, and go below the minimum number of staff required, as well.

And then, some of the long-term mitigators that are under consideration but not for immediate action, firstly, is reevaluate our KMF locations, whether there are useful alternates that would help mitigate these risks, and also if there is a potential for reconfiguring the number of TCRs that are needed, how they're geographically distributed, how their roles may overlap, and so forth.

And then, areas we're exploring right now for our updates to our DNSSEC Practice Statement is to elaborate on the triggering conditions so that, for the contingency scenarios we can envisage—obviously, there are more today than there were when that document was first written—mapping out more precisely what the triggers would be to enact the certain contingencies. I think that might be the last slide? Yes. Thank you very much. I'm happy to answer any questions.

JACQUES LATOUR: Thank you, Kim. Any questions? You can input them. Oh, sorry, Kathy.

KATHY SCHNITT: Yeah, we do have a question in the chat, a question from Steve Crocker: "Do you have statistics on how often people do not show up for key ceremonies? Do these statistics cause you to reconsider how many people are needed and/or whether travel support should be provided?"

KIM DAVIES: Thanks for the question. We don't really have much in the way of no-shows in that folks' inability to travel for key ceremonies is usually determined well in advance. We typically poll for attendance six months in advance and that is how we select the date for the ceremony, based on TCR availability and staff availability. It's relatively rare that a last-minute cancellation needs to happen.

With respect to travel support, that was definitely a topic of interest in the early years of KSK operation. We actually do provide travel support to the KSK ceremony for the trusted roles since 2014, so this is no longer an issue. Certainly, TCRs today can avail themselves of travel support if that is a particular blocker.

I think, with respect to the current health emergency that is ongoing, it's not so much a matter of funding as it is a matter of quarantines, government restrictions on travel and so forth, that are simply beyond the control of the participants to know or to predict.

So, whilst companies often have policies against inessential travel, certainly ICANN itself does in that regard. From an ICANN perspective, the key ceremonies are essential travel. From those that we've asked that are involved in the next ceremony, they've indicated that their company policies are essentially the same. So, there's no suggestion that that will necessarily be a blocker but if governments, obviously, prohibit travel, that changes the equation quite substantially.

JACQUES LATOUR: Any other questions? I guess I have one. Is your contingency to find more than one-quarter of keys material? What's your process to enable that?

KIM DAVIES: We don’t have a formal process for that so that’s something that we’re actually studying right now. We’re exploring how viable it is, what the risks are, what the compensating controls might need to be, if we’re ready to go down that path. I think, to be clear, at this stage we see no indication that we can’t do the ceremony and, in terms of the list of options, that’s pretty far down the list of things we think we want to consider. Nonetheless, I think it is important to consider the impacts and so forth to plan for it. My hope is we would never need to exercise something like that.

JACQUES LATOUR: Okay. Any other questions?

KATHY SCHNITT: At this time there are no questions in the chat.

JACQUES LATOUR: So, were you expecting more feedback from your KSK Rollover from the community?

KIM DAVIES: Not necessarily. I think that much of the feedback was from board community groups like SOs and ACs. I'm confident that the 11 comments is actually representative of a much larger number of people that have considered it.

As I mentioned, I think the only thing that was a little surprising to me was the specificity of some of the feedback. And so, we'll take that on board to study a little bit more. If anyone has any opinions on how we should interpret that advice and whether our particular approach in this instance was the right one, or maybe it could have been tweaked, it might be useful the next time around we're contemplating like this. That would be useful for me, at least.

JACQUES LATOUR: Okay.

KATHY SCHNITT: We do have one more question for Kim. "To minimize health risks, would ICANN consider arranging individual chartered planes to fly down the trustees of the key?"

KIM DAVIES: We would consider anything. I'm not sure whether it's viable or not. I think, at least, again, for the next ceremony no one has indicated that that particular concern is one that is decisive in their ability to travel. It's more a case that they're happy to travel via the normal means that they would on a normal ceremony but the risk that there is, come six weeks' time, roughly, when that ceremony is scheduled, government restrictions and so forth may be very different from what they are today. That's not a matter of the method of travel.

JACQUES LATOUR: Russ, did you have a question?

RUSS MUNDY: Thanks, Jacques. Quick question, Kim. On the review of a comment that you've seen so far, I know that there was a hope to have the next KSK Rollover sooner rather than later. Do you think that, from what you've reviewed so far in the comments coming in, that that will still be able to be the case or, as a result of the comments and inputs there, the next KSK Roll will push out to be a little while later, like a year later than what was initially envisioned?

KIM DAVIES: I think at this stage it's a little premature to speculate on the potential timeline, simply because of this coronavirus. We have a relatively small team and, truth be told, we have really needed to divert our available operational resources into holding the key ceremonies in the short term, so we've needed to suspend our forward-thinking work, at least in the interim.

Hopefully, in the future weeks, the situation will clarify somewhat and we can return to our forward-thinking project development for the longer term, but as of now there are just so many unknowns that I wouldn't really want to speculate on that front.

RUSS MUNDY: Okay. Thanks very much.

JACQUES LATOUR: Alright. Any more questions? I think we're pretty much out. Thank you, Kim, Wes, for good presentations. I guess we're on break for the next 15 minutes.

KATHY SCHNITT: We are on break. You do not have to disconnect. You are able to stay connected. This is the same room we're going to be using. Everyone, just come back at half past the hour. Susanne? You can test now, if you want.

SUZANNE WOOLF: Okay, thanks a lot. Can you hear me?

KATHY SCHNITT: I can hear you.

SUZANNE WOOLF: Awesome. How is the sound quality? Is it okay?

KATHY SCHNITT: Sounds good.

SUZANNE WOOLF: Fabulous. I am now going to see how the video will go. Hello, video.

[DAN YORK:] There you are! Yay.

SUZANNE WOOLF: Cancún it isn't but it's a borrowed office in Reston.

KATHY SCHNITT: They need some stuff in the background, there.

SUZANNE WOOLF: Yeah. Well, this is the visitor office. This is the swappable office.

[DAN YORK:] Got you. A visitor office.

SUZANNE WOOLF: Yeah. Well, I figured that sitting in my attic alone for the ICANN meeting was not something I was excited about.

[DAN YORK:] Yeah. Our comment [as a panel is] not that you have to have video. I mean, if it starts to break up or something, drop it. In the meantime, we had other comments that I was nice to see people.

SUZANNE WOOLF: Yeah. Frankly, I made sure that I was ready to do video just because that's what I do at home, too, because it helps me focus. Even if nobody cares about seeing me I care about having to look like a grownup.

[DAN YORK:] "You're wearing a shirt today," or something [to that matter, yeah].

SUZANNE WOOLF: Anyway. Jonathan is offering to show us Cancún in the background. You're the man. Thanks so much.

KATHY SCHNITT: Oh, man.

[DAN YORK:] Swap in the virtual … Actually, Zoom has … I have colleagues who have a green screen. They have it behind them and then they can put in whatever background they want. Zoom has those virtual background things so we could get a picture from Jonathan and put it in our background, there.

SUZANNE WOOLF: We can all be there together with him.

[DAN YORK:] Picture this, if you were.

SUZANNE WOOLF: You're bringing the beer too, right?

[DAN YORK:] You could be drinking the beer, yes. You would never know. That could be in your coffee mug right there for all we know.

SUZANNE WOOLF: No, it's coffee.

[DAN YORK:] On that note, I'm going to drop off and go get some tea.

SUZANNE WOOLF: Yeah, likewise.

KATHY SCHNITT: Vittorio, do you want to go ahead and test your audio?

VITTORIO BERTOLA: Yeah. Can you hear me?

KATHY SCHNITT: I can, thank you.

VITTORIO BERTOLA: Yeah, hi. I think we can also try the video. Hi.

SUZANNE WOOLF: Looking good, sounding good. It's fine.

VITTORIO BERTOLA: Yeah, okay. Because of course, you know, everybody is working from home in Italy now so all our Internet access networks are [signally] strained. And so, you never know.

SUZANNE WOOLF: It sounds like a really good idea to not get sick in Italy right now.

VITTORIO BERTOLA: Yeah. I mean, I hope it doesn't get like this everywhere. Unfortunately, the rest of Europe could be like this in a week or so. We don't know.

[DAN YORK:] So, I was offering to show you Cancún. Here it is, if you want.

KATHY SCHNITT: Yes, beautiful.

SUZANNE WOOLF: I'm going to go shuffle down the hallway in the office to get another cup of coffee. I'll be right back.

KATHY SCHNITT: Vittorio, do you want me to manage the slides for you.

VITTORIO BERTOLA: Yeah, maybe it's better that way.

KATHY SCHNITT: That is just fine. You just say "next slide" whenever you're ready.

FRED BAKER: Kathy, so we have our Chinese colleagues on the line yet?

KATHY SCHNITT: Let me go see if anyone has joined. Nope. Neither of them have joined yet. We will keep an eye out for them, Fred, and then we'll transfer them to panelists as soon as they do join.

FRED BAKER: Okay.

JACQUES LATOUR: Yeah. I see their [poll and audio showing] I guess.

KATHY SCHNITT: That looks very nice compared to where we are, huh, Jacques?

JACQUES LATOUR: Yeah.

[DAN YORK:] I had the opportunity to stay in Seattle or to come to Cancún, so I think I made the right choice of corona.

JACQUES LATOUR: Bravo.

[DAN YORK:] Right here behind me is a dolphin – they have dolphin shows. There are six dolphins [inaudible]. You can see the beautiful ocean shores, here. It's a lovely city. It'll be great when we come back here in a year. The people are very nice here. There are a lot of amenities local and around the conference center. So, it'll be a really good location.

JACQUES LATOUR: Good. Thanks for the quick peak, there. Thanks.

UNIDENTIFIED MALE: Kathy, are you able to kick him out of the room?

UNIDENTIFIED MALE: Yeah, can you please turn your video off? I don't want to see that anymore.

KATHY SCHNITT: I can do that.

[DAN YORK:] I comply.

KATHY SCHNITT: No, Jacques, it's everyone.

JACQUES LATOUR: Okay.

KATHY SCHNITT: Okay, Fred. We're right at the start. Do we want to get going?

FRED BAKER: Sure.

KATHY SCHNITT: We'll start the recording. Welcome to the ICANN67 Virtual DNSSEC and Security Workshop on Wednesday the 11th of March 2020, part two.

Today's meeting is being recorded. For the panelists, please remember to state your name before speaking and please keep your phones and microphones on mute when not speaking to avoid any background noise.

As we are utilizing the webinar rooms for this virtual session, attendees will only be allowed to ask their question and/or comments through the Q&A or chat box and we will be happy to read those questions or comments aloud. This session will run until 17:15 UTC. I will now turn it over to Fred Baker. Please begin.

FRED BAKER: Hi, there. So, this particular session is trying to look at the security in practice of DoH servers and the DoH protocol. So, we have three speakers that I hope will be interesting. One is Vittorio Bertola, who will be looking at the resolver discovery problem and [inaudible], along with [inaudible] Security and Stability Advisory Committee report on DoH. And then, too, from China, GenShen and Gong Yiming. They found an issue with malware that is being used by DoH services to control a botnet. And so, I'd like them to talk about their discovery and what the implications are. In any event, that's the conversation that we have this morning. Vittorio, can I turn this over to you?

VITTORIO BERTOLA: Yes, sure. Good afternoon, everyone. My presentation is about the problem that is currently being discussed of how to discover DoHs on the local network. Next slide, please.

Okay. So, basically, traditionally, if you are a traditional DNS clientele, such as the [inaudible] of an operating system, you have a very simple algorithm to choose which resolver to use. Basically, the question that you as a [inaudible] ask yourself is, did the user enter a specific resolver? So, if the user had chosen to use a specific resolver, of course, that's the one you will use. Otherwise, normally what happens is that you've just asked the network for a local resolver, which is usually advertised to you by the local network via HTTP or equivalent IPv6 things.

Unfortunately, this mechanism doesn't work if you want to use a DoH because, even if you can get to the IP address of the local resolver through this mechanism, and even if you assume that the DoH resolver is also the local traditional DNS resolver, you need the URI template to contact the DoH resolver. Having the IP address is not enough. Next slide.

So, basically, this is what is evolving, at least in this first year or so of DoH implementation. So, of course, every application will just allow users to enter a specific resolver and, of course, if they choose a resolver then that's the one that will be used, but very few users actually choose a specific resolver.

So, in the end, there is a new question that applications are asking themselves, which is, "Do I have a preferred DoH resolver?" This is what Firefox have decided to do. So, in any case, there is no choice [inaudible] by the user. They have their own preferred resolvers and they will use those ones, at least in the US. These often have several issues which are out of the scope of this presentation.

So, let's say I would rather stick to the traditional way of doing this and just ask at the local network for a DoH resolver. But this is currently impossible because there is no way, currently, to ask the local network for a DoH URI template. So basically, what the other applications like Google Chrome or Windows are doing is that they are trying to build a database of, basically, which DoH resolver is available over which network so that they can then pick the resolver from that list. Next slide.

So, how do you do this? Well, it's not just that you don't have a way to ask the local network for a DoH resolver. Also, you have, really, no way to know which ISP you are on, so there is no way to ask the local network who your ISP is. Even if you had this, even if you knew which ISP is serving the customer or the user of your application then you will still have other problems.

One is that several people say/point out that a DHCP is not authenticated so it's not a very secure way of asking anything for the local network. The other problem is that even users that have not configured a DoH resolver would have configured an original DNS resolver different from the one that you get from the network.

So, in the end, this is why Google, and then also Microsoft, have devised this temporary solution, which is what they are implementing for the first DoH releases, which is where, basically, they look at the IP address of the traditional DNS resolver and create the resolver which is configured in the operating system, which usually could be, at least the one that comes from the local network, but could also be something that the user has entered.

And then, they will use that IP address to, basically, understand which operator the user wants to get the DNS from. So they, basically, have a hard-coded conversion table in which they have the public IP address of the resolver, and then they also have the URI template of the DoH resolver, which is provided by the same operator.

So, at least for the ISPs that have already started to provide DoH resolvers, this allows, basically, the application to just start using the correct DoH resolver that matches the DNS resolver that has been entered in the configuration.

This works, also, if the user is a non-local DNS resolver. So, it’s a good way to approach the problem, but then—next slide—there are problems with this solution.

The first one is that you need a hard-coded list of, basically, each and every DNS operator. Of course, this can be done on a very small scale as a temporary solution but, I the end, you cannot imagine that someone can create a list of each and every ISP that has a DNS resolver and wants to make it available via DoH.

And so, this is problematic in the long-term because it leads to further centralization. Also, it doesn't work if the resolver is a private IP address because, of course, you cannot use it as the key in this conversion table. This is the common case for access networks by Telcos, at least in Europe. And because in these networks, usually, the home router is asking as the user’s DNS resolver and the acting as a follower to the primary DNS resolver in the platform of the ISP.

In some cases, even the main resolver platform of the ISP or of the network, especially in corporate networks, actually has a private IP address because, for security reasons, the operator doesn't want to make the resolver used for its customers available to the general public. And so, this is also a case which doesn't work with the solution. Next slide.

For this, it’s the difference between what the browsers that have been trying this way of doing things are expecting and what ISPs are actually doing. Meaning, in their model it is envisaged by the browsers, basically, the computer/the user device would ask the home router for a resolver and it would get the public IP address of the main resolver platform.

But, unfortunately, what happens in reality is that the home router is configured so that it gives its own IP address, which is, of course, a private IP address to the user. And then, the user device uses that home router as a DNS resolver and the home router is forwarding the queries to the main resolver.

And so, basically, this is the problem we are faced with. ISPs that still want to provide DoHs over to the user and have it supported in Chrome and in other applications cannot do so because there is no way of doing this with this current method. And so, the discussion is ongoing with the browsers on, basically, how can we find a solution for this? Next slide.

So, there is already a draft for resolver discovery, which, by the way, is authored by Paul Hoffman. The draft tests two different methods for asking the local resolver and the local network for, basically, information which is general information but could be the URI template of the DoH version of the same resolver.

And so, one of the methods is via DNS through a specialized resource type for the reverse IP address, and the other one is via HTTPS. So, [let’s say a] document for a well-known URL. Both methods have problems.

First of all, there is a problem with the fact that you have two methods which require people to implement two different things, otherwise you risk them not being compatible with everyone. The web-based method has big problems because, still, it doesn't work in the scenario we saw because it connects to the home router, not with the main resolver platform.

And also, the most important thing is that it requires heavily modifying all CPEs/all home routers because they don’t currently have a website or, at least, they have—and sometimes they have—they don’t support, anyway, this kind of request.

The DNS method works better in this scenario because it just requires adding support for this query on the main resolver platform. But again, the current method doesn't work with private IP addresses. And so, at least it’s not easy to make it work with a private IP address. Next slide.

So, in the end, this is the idea that is being considered, basically, to use the DNS method but make it work also with private IP addresses. And so, the idea could be to use a special use domain name so that the user’s device/the application can just make the query, and the query can also get forwarded throughout the chain of forwarders, and then it will get the reply from the main resolver, and the [inaudible] will be, basically, usable.

But at the same time, to address the issue of security, there is the need for a sort of second-factor authentication of the network. You ask the network, you get the template for the DoH resolver, but then you need to have another way to verify that [inaudible] makes sense, and so to, at least, counter attempts to attack the user by giving a different resolver.

This could be done, again, through a list, through a form of hard-coded list as a second-step temporary solution, or maybe through checking some certificates or other ideas. This is a subject for discussion.

So this, still, is not useful for those applications that by principle decided they don’t want to trust the local network and the local resolver, so it doesn't address, let’s say, the Mozilla deployment model, but it would be a reasonably secure solution for the applications that actually want to trust the local resolver and, basically, don't want to change the resolver that is used in the system from the one that is used for traditional DNS.

So, I guess this will be one of the discussions that will be had in the newly created [ADD] Working Group or the IETS. This was also meant to spread the word about the discussion and encourage people to participate in this new working group. With this, I think I'm done so I will be happy to take questions.

KATHY SCHNITT: Vittorio, we do have a question from Daniel Migault: "Are there any other fields that need to be discovered other than the authority part of the DoH URI?"

VITTORIO BERTOLA: Well, I think this is also open for discussion. Basically, if you just want to establish the DoH connection and use the resolver, I think the URI is all that you need. But it might be useful to get other information, for example, to do this kind of two-factor authentication or verification of the actual identity of the resolver it's on. So, I think it's up for discussion.

KATHY SCHNITT: And at this time, Vittorio, there are no further questions.

VITTORIO BERTOLA: Thank you.

KATHY SCHNITT: All right. We'll hand it back over to you.

FRED BAKER: Okay. Well, thank you, Vittorio. Suze, can I get you to talk, now?

SUZANNE WOOLF: Here I am. How is the sound? Can you hear me now?

FRED BAKER: I can hear you fine, yeah.

SUZANNE WOOLF: Fabulous. Okay. I will leave video on unless it seems to be causing problems with the audio. Thanks for the presentation so far. It's always a good thing. What's I'm going to be presenting is a quick overview of a new SSAC paper that's about to come out on the implications of DNS over HTTPS and DNS over TLS.

Vittorio just gave us a very detailed, coherent, and cogent discussion of a particular aspect of the challenges with deploying DoH. What the SSAC paper I'll be talking about does is it steps back to a more abstract level.

SSAC was very concerned about the implications of DoH and DoT from the perspective of the ICANN community participants in the DNS ecosystem. So, what we worked on was presenting an overview and a high-level view of what the implications were on more of an architectural level than an operational one. Although we do discuss the technical details it's toward attempting to derive a kind of higher-level view. So, next slide, Kathy, if you would.

One way we started trying to figure out how to say something useful to the community about this—Barry Leiba and I, co-chair of the SSAC Work Party on this—as soon as we got into it we discovered that it was going to be very complicated because there is a lot of technical detail, there is a lot of nuance, and there is a lot of complexity because we're changing a number of things at once about how DNS functions in context with other network services and with applications.

But the outline on the paper is we did an explanation and comparison of DNS over HTTPS and DNS over TLS, focusing on standardization and deployment status so that people can see what's coming.

We kept asking questions of each other about the implications and hearing, "It depends." So, we ended up writing up an exploration of the effects and perspectives of several different groups of stakeholders, parents and other people with responsibility for others: enterprise network managers, dissidents/protesters, civil society folks, and Internet service providers, which is by no means an exhaustive list but let us explore multiple perspectives.

It discusses examination of application resolver choice and what implications come from these decisions, which, frankly, DoH and DoT put the choice of resolver largely under control of the application in a way it wasn't before. We have a brief discussion of implications on the namespace due to DNS stub resolution moving to applications. Next slide, please.

We did find it was an unusual undertaking from SSAC in that there are several things SSAC often does that we're not doing in this paper. We discovered we couldn't agree on right and wrong labels with respect to DoH and DoT, their implementation and deployment choices.

We found we had to stay away from "more privacy is always better" or "more encryption is always better" because, again, we ended up exploring different perspectives because it's very hard to make flat statements here.

Because of that and because we were looking at the ecosystem overall, including many players who are not ICANN participants, are not part of the domain name industry, and so on. We ended up not be able to make specific recommendations to the ICANN Board. So, the goal here is to inform the community and to expand people's perspectives. Next slide, please.

The conclusions we did discuss is that evaluations of DoH or DoT rely on the perspective of the evaluator, how they're implemented, how they're deployed, how settings are configured, and who uses them.

Regardless of perspective, we conclude that the deployment of DoT and DoH will be disruptive, mainly in the implementation and deployment of the technology because they change the control perimeter around the network or the application or the user.

Application-specific DNS resolution through DoH and DoT does present a host of legitimate challenges: how networks and end-point interact, who has access to DNS query data, how to protect and manage networks in this new model. We need to go to slide 15, please. Thank you very much, Kathy. This was a longer presentation and I'll ask that it be posted so people can review because there is a lot of technical detail on the part I'm skipping over in the interest of time. Next slide, please.

One of the stakeholders we looked at … It's easy to call it "parents." It's, frankly, anybody in a position of responsibility for others but, particularly in the western context, the paradigm is around parents who might wish to control children's access to the Internet. DNS can be an effective control point for this, which turns into something of a proxy for DNSes' control point for a number of mechanisms for controlling network access.

Services have always existed to provide this type of blocking and children and anyone else have often been skilled enough to work around them. It does seem that DoH particularly will make this kind of blocking more difficult. Next, please.

We looked from the perspective of network managers. Many different types of organizations can be considered enterprise networks. Basically, you have a controlled perimeter around particular activities or resources.

Those managers often have a positive obligation to understand and control the traffic on their networks and DNS is an important control point for them, too, and the introduction of new DNS transports threatens their control, also – control on management.

Regardless of the purposes to which it is deployed, the challenge is that the model changes. Who sees what information, who has control over what information changes as these new technologies are deployed. Next, please.

For folks who might have a need for privacy—which may or may not be, frankly, related to their interactions with others around them or with authorities—the Internet is an important vehicle for dissidents and protestors to spread alternative views, critique politics, shed light on corruption, human rights abuses. It's a powerful force for diversity of views.

By encrypting DNS queries and resolution, DoH and DoT can help [shewn] users from being tracked by their ISPs or governments. There have always been personal and private access to VPN resources but, like other technologies in the space, DoH or DoT can help, can promote privacy, and promote diversity of access and views, but this is largely a matter of policy and greatly influenced by governments.

There is no fantasy here and DoH and DoT are not going to be where that changes. Next, please.

ISPs. A lot of the motivation for deploying DoH and DoT is sometimes described in terms of business practices, operational practices of ISPs, that some users don't care for. They collect data about users that they exercise control over users' access to resources. In particular, we found ourselves discussing many governments obligate ISPs to block traffic using DNS as a control point. DoH and DoT may mean that ISPs now become obligated to block traffic using other means.

Some ISPs may resort to blocking all DoT traffic or offer their own DoH or DoT services. This is where the work that Vittorio was discussing is extremely important. The discovery problem is the difference between having your application provider choose your resolver for you and having some control over what choice is made within the application for resolver use and transport use. ISPs may blacklist known DoH servers based on known IP addresses but this will not work 100% for the reasons that Vittorio discussed.

So, again, lots of complexity. A lot depends on the implementation of specifications and standards but a great deal depends on implementation choices and configuration choices. So, there are multiple dimensions here to how these technologies work in practice. Next, please.

So, the final thing that we discovered that we were interested in discussing, and where we hope to say something that is particularly of use and interest to the ICANN community and to those in the domain name business—next slide, please—is that we found that changes in access to the namespace, changes in access to DNS query resolution—both by handing the resolver choice to the application and by encrypting the channel for queries—has some potential implications for the DNS namespace, and how it is used, and how it is accessed.

Applications performing DNS functions themselves may cause other disruptions, which is not the direction that users are used to looking for behavior on the network. One industry concern with respect to applications providing DNS is that they will undermine the usefulness of DNS as a generic and protocol-neutral naming system for the Internet.

What that means is just that the more DNS is deployed and managed as, frankly, a component of the web, which DoH particularly promotes, the more concern it seems legitimate to have that DNS will be increasingly optimized for the web, which it already is in many ways, and that it will be harder to write other protocols that don't depend on the web but rely on DNS as a generic and protocol-neutral naming system for the Internet.

So, we believe quite strongly that that's something that needs to be considered and the folks in the industry and the DNS ecosystem need to be aware of. Again, namespaces may become tailored to the requirements of a particular application of protocol.

Web browsers have begun to cache web content per origin. In practice, this means each browser tab, if you use a browser directly, now has its own cached versions of content. We are concerned about the tendency to do more of that.

I think that's the end of the slides. We can have discussion/questions. We had actually hoped to have the paper out by now but it turned out to be challenging to make sure that we got to a real SSAC consensus. So, that will be appearing in the next day or so. Thanks.

FRED BAKER: Well, thank you. Kathy, do we have any questions or comments to raise?

KATHY SCHNITT: Fred, not at this time.

SUZANNE WOOLF: Maybe we'll see some after the last presentation, then. Thank you. I am going to shut off video in the interest of bandwidth conservation.

KATHY SCHNITT: Thank you, Suzanne.

FRED BAKER: Okay. So, next question: do we have our Chinese colleagues online? They weren't at the time the show started.

KATHY SCHNITT: Yes.

FRED BAKER: Ah, GenShen Ye. I see you there. Okay. GenShen, can I turn it over to you, then?

GENSHEN YE: Okay. Thanks, Fred. Hello, everyone. I'm GenShen, with my co-worker, [inaudible]. [Reference list XP] on the Network Security Research Lab. Our team focuses mainly on [two alias] and on other related research and DNS data security.

Today, I'm going to talk about the DoH behavior [where above] [inaudible] in the [inaudible] we discovered in mid-April last year. At [our work], we have a pretty good Internet [inaudible] monitoring system going which detects new and interesting threats on a large scale. On April 24th last year, the system for [inaudible] because the [inaudible] [loaded up as a symbol, as a magical number of God].

[inaudible] and, as we know, north of the [inaudible] have [two continents], the [inaudible], AKA "commander," and the [Cantaro] and, the supports. The supports are [inaudible] which connected to the C2 to gather instructions on what supports [are needed to].

[In other words, for the] industry, a common practice to deal with [inaudible] is to deal with the C2 because [it is difficult to clear all the boards] where, at the same time, the C2 can be cleared or be [broke]. Supports become [headerless] and the [center] is [inaudible], the reason we find the [inaudible] interesting.

Our team has been [trading] [inaudible] for the last six years and the [inaudible] is one of the most sophisticated [inaudible] we have ever seen regarding hitting its C2s. First, it is the most common technique hard-code [inaudible] code. Nothing fancy here. [inaudible] when [inaudible] .com is regularly used for this service. So, it is certain that we are not [inaudible] any security products.

One way [inaudible] is to access specific [inaudible] URL, [inaudible] [contender] and on the page, to gather the real C2 information. Besides [inaudible] being, [inaudible] [also used to have, too] [inaudible] information. Same logic with [inaudible]. So, using [inaudible] is pretty fancy but [inaudible] yet. [inaudible] also uses DNS over HTTPS. [That] is something we have never seen before. Okay. Next slide, please.

The background shows the [inaudible] of the [inaudible] redundant mechanism. When [inaudible] ports go online they need to go through three stages to get a hold of the final C2. Each stage breaks down to two steps. I'm going to go over each stage in detail in the following slides. Next slide, please.

Stage one. The [border] has three choices to start with, here. You can use as a hard-coded [cyber] text or [inaudible]. Let's take a look at the [inaudible] one as an example. [In the square] [inaudible] you can see when you access [inaudible] URL, a page with [an incredible messaging] is loaded. By the way, [inaudible]. You can [inaudible] URL and you can gather the text here.

After the [inaudible] message, the [inaudible] a URL shows up. It is [inaudible] configured.com. So, let's press and let's start a PNG. So, domain name looks like a [legit] but it is actually a malicious domain controlled by the [inaudible]. [inaudible] the above URL as input and as a string starter. So, the URL changes to HTTPS, the domain, and as a [inaudible] the PNG.

Then, the [border] accesses a URL to download as a standard .PNG file. The file is actually a [RUA] file. One is the [inaudible]. Another domain name appears, so it is [t.crawlerapb] configured.com. Notice, here, the border only knows about its domain name, not its IP address.

With that, it goes to stage two. Next slide, please. Stage two, the [product], again, has two choices here. They can [figure] as a [inaudible] or DoH to [harbor it as a tagged file, one] .PNG. Here, we can take a look at the screenshot. It is [inaudible] request was sent to the [inaudible] DNS over HTTPS resolver.

We asked for the [inaudible] [resolver] for [cheap.cloudabconfigured.com], which we [inaudible] stage one, and the [Cloud's first] DoH resolver returned [a text recorder], which is [encrypted].

After [decryption] we can see the [real text column] just here. [inaudible] No, it is [imageone.crawlerapbconfig]. [It's about that is above you] as input and as a string run. So, the end URL changed to HTTPS, the domain, and, as a slash, run.png. Then, the [border] downloads the tagged file run.png from the URL. [inaudible] [see it] and the final command and the [inaudible] domain name shows up. So, it is [c.crawlerapbconfigure.com] again.

Notice here the [border] only knows about its domain name, not as an IP address. [Instead], it goes to stage three, the final stage. Next slide, please.

In this final stage, [inaudible] asks, "[Cloud] failed DoH resolver for the [inaudible] C2 domain name [c.cloudapbconfigured.com]." In the [inaudible] you can see "DoH requester is [sent] to [cloud failed] DoH resolver" and a record of 43.224.225.220 is returned. With the final set of information, the [inaudible] with the C2 and the [beginnings, they are conversations no.] Next slide, please.

This slide will take us [inaudible], in which we can see that [inaudible] DNS [recorder section] with the performance of the DoH function, and that is [really easy]. Next slide, please.

So, we have quickly gone through the three stages [inaudible] need to take to communicate to their C2s. What I have covered is just some [inaudible] of the [lands]. Folks are welcome to check out our original blog, which has all the details and everything [inaudible]. [inaudible] [researcher], where, basically, [inaudible] of the C2s from a DNS perspective.

DNS monitoring and blocking is a very cost-effective and easily scaled security weapon that defenders can use. [They are giving] the security providers and the products [inaudible]. For example, our team is behind one of the biggest DNS service providers in China. Every day, we block about 20,000 active malicious domains. [Where we will] not be able to see or block the DNS requests at all, if all the [boards] started communications with using DNS. [inaudible] is the first [inaudible] [we've seen] using DoH for either one of these last one to [inaudible]. So, there are some things that community either to think about [inaudible]. Okay, that's all. Thank you. Does anybody have questions?

KATHY SCHNITT: Thank you, GenShen. Fred, back over to you.

FRED BAKER: Hi, there. So, Zoom is being very strange right now but I want to thank each of the panelists. I think as far as time goes we're four minutes before the top of the hour or the top of the session. Kathy, so we have any questions?

KATHY SCHNITT: At this time, we do have no questions.

FRED BAKER: We have no questions. Okay. So, we have two choices. Wes mentioned that he had a slide that he might want to share and talk about. Is that something you can do in three minutes?

WES HARDAKER: Yeah, it's a quick slide.

FRED BAKER: Oh, okay. Well, why don't you talk about that, then?

WES HARDAKER: Sure. Let me share it really quick. Basically, as many people know, Firefox has rolled out in the US and one of the things that they offered when they were turning this on by default was that ISPs could stand up a domain and, if they get an NX domain answer for it, they will not turn on DoHs. They will not send their Firefox browser resolver request to Cloudflare. They did that trying to figure out … There are some ISPs that really don't want to let users use DoH. Users can then still go force it on anyway but it gave the ISPs an option to turn off that by default.

I have been measuring that usage through about 1,000 RIPE Atlas probes since October of last year. This graph that I am … I am displaying a graph, right? Everybody can see the graph?

KATHY SCHNITT: Yes, Wes. We can.

WES HARDAKER: Thanks. So, this graph has been measuring that. You can see that it has climbed over time. It started at 4.5% back in October of last year. It goes up and down a little bit. Some of the noise comes from the number of Atlas nodes that are actually participating or go on and offline.

And then, recently, actually, and it looks like near February 25th of this year, it took another big jump upward. So, we're now up to about 10% of the RIPE Atlas resolvers that are contributing data to this now actually seeing useapplicationdns.net as being spoofed with an NX domain saying it doesn't exist when it really does in order to convince Firefox users to not use DoH.

So, just a point of order/another data point for how to measure the popularity of DoH. This in no way affects the good presentation that just happened about DoH being used by malicious software because they are not going to check and use this domain at all. Any questions? All right, good.

FRED BAKER: Okay, [failing that]. Thank you very much, Wes. That's interesting. Kathy, I think I need to turn this back to you, correct?

KATHY SCHNITT: Thank you. At this time, we're going to go into our next session, DNSSEC provisioning with third-party DNS providers. I'm going to hand this over to Steve Crocker.

STEVE CROCKER: Thank you very much. While I get organized here … Oops, sorry. I said I was adept at this but I turned out not to be so adept. Excuse me. We're going to have six panelists. I think five I have seen online. I'm waiting for one more. Kathy, I hope you were able to promote the two of them to panelist level.

KATHY SCHNITT: Yep. Gavin and ...

STEVE CROCKER: Brian Dickson.

KATHY SCHNITT: Yep, they're promoted.

STEVE CROCKER: Good. I am trying to—I promised that I was good at this but I turn out to be fumbling here—get the appropriate thing up here. There we go. Okay. Now, I have to hit "share," and then I have to hit "desktop two," and then I have to hit … Okay. I believe I've got it. Everybody can see "DNSSEC provisioning with third-party DNS providers," I hope?

KATHY SCHNITT: We can.

FRED BAKER: Yes, we can, Steve.

STEVE CROCKER: Great. Thank you very much. All right. So, if you noticed in the program, this was heralded with a yellow highlight indicating this was at the expert level, so we'll see how this goes.

When DNSSEC was defined, one detail is that every so often, the child zone, the key for it would be changed and there would have to be an update to the DS record in the registry.

As it turns out, there is no well-defined path for this to happen. The natural path for updating things is from registrar to the registry and the registrar does not have a natural interface commonly defined for third-party DNS providers.

So, this has turned out to be a bit of a stumbling block. This is one of two DNSSEC provisioning details that have gotten my attention. The other one is, what happens if you have multiple third-party DNS providers, each of which is using its own keys to sign the zone? How do they coordinate across that?

These two issues are ones that I'd like to pursue over time or see others pursue over time. In this session, we'll focus just on the first one, how do you update the DS record when you're using a third-party DNS provider that is using its own keys to sign the zone?

Implicit in what I'm saying is, how do you do this in an automated way? You can certainly do it in a manual way if you, as the registrant, sit there and interact with the DNS provider, copy at length, carefully, a new key, and then go into the web interface for the registrar and try to type that in. And so, it could push a new DS record upward.

That's the verbal summary. We'll go through all that a little more carefully. We have, today, six panelists chosen from across registries, registrars, and DNS providers: Jim Galvin, Irwin Lansing, Gavin Brown, Brian Dickson, Jothan Frakes, and Ólafur Guðmundsson. I hope they're all here, actually, but we'll see in a minute.

I'm going to talk for a few minutes just to frame the question more precisely, and then I ask each of these people in the order listed to describe what they are doing and how they'd like to see things move forward.

Then, I'll go through a second round of asking the panelists to comment on what they've heard or anything else they want to add. And then, we'll open it up for questions. Then, I'll close up the session with a few comments. So, that's the order of events for the moment.

So, as I said, today's focus is the DS update problem. Third parties generate keys, sign the zone – there is no well-defined path. There has been some progress and one of the paths of progress is documented in RFC 8078, as I see at the bottom of the slide. The other DNSSEC provision problem of coordinating multiple third-party DNS providers we'll take up in another venue.

So, this picture is intended to show the components of a process. So, you have a master copy of the zone. That is updated by the registrant adding or deleting new records as needed. And then, those are passed down to the DNS provider who does his own key generation, signing, and publication of the zone in his own name servers.

And the missing piece here is, how does a new DS record get up to the registry? There are three different dimensions that describe what I think are the entire space of possibilities. One is whether or not the DS record is pushed upward by an action of the DNS provider versus whether it's pulled up – that is, it's published by the DNS provider and then the registry or one of the other entities pulls the zone and signs it and then copies it up.

The second dimension is whether or not it's the registry, or perhaps the registrar, or even the registrant that is at the other end of that process. The third dimension is the choice of whether or not it's the DS record or the KSK or both. And so, there is a three-dimensional space of possibilities.

In the picture that describes what I call the "pull model," or by pulling, the DNS provider publishes the new record and publishes it in the form of either a CDS or a CDNS key or both in the child zone. And then, that's noticed and pulled upward either by the registry, which is the solid line, or … I don't think we have any work examples but it could be done by the registrar or it could be done by the registrant, either manually or in some automated fashion if he’s running some auxiliary software.

The alternative is that as the zone is signed there is an active effort to push the new DS record upward. I don't know of any—I stand to be corrected on this—active methods that are in process, that people actually use this. Although, Domain Connect has been envisioned as a possible path for allowing new DS records to be pushed to the registrar or to the registry.

So, on today’s panel, we have people who are well-placed throughout different parts of the ecosystem and have asked them to speak to the current state of automation. One element of this is whether or not it’s okay to bypass the registrar. That is, I’ve heard, in some cases, registrars say, “Well, we don’t like the idea that there is a separate path to the registry because then we would not know all of the records on behalf of our customer,” but I’ve also heard the alternative, which is, “It probably doesn't matter for the DS records.” So, that’s one small point of discussion.

There has been some forward progress but not quite enough. I suspect that in the gTLD environment, that is the parties that are contracted to ICANN, there is no well-ordered path and no forward progress, as best as I could tell.

So, I asked the question, “What are the next steps for complete automation and what are the impediments?” So, here is the set of panelists, again. With that, I would like to go through, in the order that's here, each of the panelists has provided me a couple of slides, different amounts of information from each one. In the interest of time, let me ask each of you to be fairly terse and pointed. With that, Jim, are you ready? I hope.

JAMES GALVIN: Yes, I believe. Can you hear me okay?

STEVE CROCKER: Yep.

JAMES GALVIN: Excellent. So, I only have the one slide to display. I'll comment briefly about it and then I'll respond to Steve's question with just a couple of points.

This is showing Afilias as a registry service provider. We have over 200 TLDs. I only call out two TLDs of interest, here. You'll see that "org," in particular, shows a 1.25% penetration of delegated signed zones in its entire TLD, and "info" has 1.18%, is what that number actually is, there. It's just a little over 1% and just under what .org has.

The label "all" there is an average across all of the 200-plus TLDs that we have, and a little over 1% penetration. That actually ranges from the spectrum of TLDs that just have a couple to we have a couple of really small TLDs with just a few hundred registrations with 97%

penetrations. You're going to hear some really good percentages from others who are going to talk here, too.

And then, the "others" category is if we take "organization" and "info" out, that "others" there at the bottom is just averaging over all the rest of our TLDs, and that's at, I think, 0.68% as I recall, penetration of DNSSEC.

In terms of what we're doing and what Afilias does across all of our TLDs, we have a long history with DNSSEC. We were the largest gTLD to sign a TLD, back in 2009. We weren't the first gTLD, in fairness; .museum gets that. That's Paul Vixie. But at the time, .org was first, followed closely by .info. Both of those were over five million domains under management at the time. This was even before the root was signed, at that time.

We then took to testing with registrars to determine what it really meant for a registrar to support DNSSEC and how that worked with the registry. In fact, that identified above in the EPP extension or DNSSEC provisioning that resulted in that standard being updated right away. And so, that was a good thing for the community, in terms of just the interaction between a registry and a registrar.

What we actually do today is we support registrars by only looking for the DS record itself. We don't look for the key record. That's what we expect a registrar to give to us. We do that because, that way, we don't have any responsibility for the hashing algorithm and the choice that's used by the registrar as a general service provider, rather than imposing those kinds of restrictions on a registrar.

And equally important on a registry that might have a particular preference toward registrars, we wanted to be able to make sure we supported the broad spectrum of things and not be tied to that. I'm sure that we'll get to others who will talk about the fact that, in this interaction between registries and registrars, you can either take a key or take a DS record. We prefer to take the DS record and that's what we do across the board with all of our TLDs.

As Steve commented on, one of the issues that I have often brought up in the past is we really do depend on our registrars to have that relationship with the registrant. So, with the solutions that Steve was talking about, when you have this option of the registry somehow being a part of pulling this DS information or providing a mechanism for DNS service providers to push that into the registry directly, we don't do that because we believe, and our understanding with/our relationship with our registrars is that kind of imposes on their standing in the ecosystem.

So, that's just an unfortunate side effect of what is allowed by the policies as understood in the ecosystem in which we work, primarily. And that is the sum of my comments, Steve. Thanks.

STEVE CROCKER: Thank you very much. We'll move onto Irwin Lansing. Irwin asked me to cover his material because he is a bit under the weather at the moment, although he assures us he’s not at risk. Irwin is also online and may offer up a comment, without speaking perhaps.

So, Irwin is from .dk, and they have an interesting situation. All of the name servers that people use have to be registered with the registry, and that means that they have quite a bit more information about those name servers and can have direct relationships. They have to be approved before you can put your name or your zone into those name servers. And then, that means that the operator can manage the [Glu], etc.

So, in order to delegate a domain, this relationship to the domain and the name server operator gives a name serve operator the option of moving things around internally but, more particularly for this context, they can manage the DS records.

In the future, they plan to move toward RFC 8078, which is the pull model with publication of CDS and CDNS key records. I hope I haven't mangled what you wanted to say, Irwin.

IRWIN LANSING: No, excellent. That’s the gist of it. The point is [the thing on top] because we have the registration of the name server by the name server operator. We can also turn that around and, when the name server operator logs in, we know which name servers he has by that we know which domains he has under management. And then, we can allow some privileges to be delegated to him in our portals and EVP, which could be bending DS, yes.

STEVE CROCKER: Good. Thank you. All right. Let me move on to Gavin Brown from CentralNic, who is going to be speaking, I believe, about Slovakia in particular, .sk. You're muted.

GAVIN BROWN: Good morning, good afternoon, and good evening to everyone. Yeah. So, I'm here to talk about CDS scanning and what we're doing with .sk, which is one of our ccTLDs and what we are looking at thinking about doing for our gTLD platform. So, can I control slides?

STEVE CROCKER: You tell me it's the next slide.

GAVIN BROWN: Yeah, if we can move on? So, this is a graph showing how DNSSEC adoption has taken of in .sk. We started signing the zone about halfway through 2019, so it has been getting on for nine months, now. We saw a pretty reasonable rate of adoption. So, we were heading toward the 3% mark until we enabled CDS scanning at the end of last year.

As you can see, we had a pretty rapid ramping up of DNSSEC adoption, to the point where we are approaching [35%] adoption of DNSSEC on the .sk domain names, which obviously is something we're very pleased with and proud of. If you can move onto the next slide?

So, we were lucky that our friends and neighbors in the Czech Republic has already warmed up our registrars. Obviously, there is a strong overlap between the registrar population in .sk and in .cz. So, CZ-NIC had already deployed CDS scanning, and so our registrars were ready for it.

We also had a lot of useful information and advice from our friends, from [inaudible] and others at nic.cz, who pointed us toward what they were doing in FRED, which obviously is an opensource registry system. There is a CDS scanning tool in there that was the inspiration for the tooling that we built for SK-NIC. We also worked on the scanning policy, which, again, was sort of inspired by one of our neighbors from switch.ch.

So, the policy we have in place in .sk is that domain names … We do a CDS query every day to see what’s there. If that CDS record set differs from the DS record set that we have—whether there are CDS records present and we don’t have DS records or the CDS records differ from the existing DS records—then the domain goes into what’s called a “pending state,” where we start scanning it much more frequently, every three hours.

But we don’t implement the changes to the registry database until there has been a stable period of three days. So, the CDS record set doesn't change over that three-day period.

We do our probing using multiple independent resolvers that we communicate with securely. We use DoH for that, as it happens. We talked to some recursive resolvers over a secure protocol to get the CDS records from multiple different vantage points. Again, all three have to agree that the CDS record set is the same in order for us to implement the changes. So, if you could move onto the next slide again?

We would like to be able to turn CDS scanning on in our gTLDs. Here are some of the challenges that we need to resolve. Some of them are technical, some of them are policy, some of them are business.

The first one is it's not entirely, unambiguously clear that we would be allowed to do it. Some people say the wording of a registry agreement says, "Well, actually, you know, it's best practice. The RFC is best practice, therefore you could propose [a standard.] Maybe you could do it without having to get ICANN's permission."

ICANN's position is, "Get an RSEP," which is a challenge because then we have to go out to every single one of our registry clients and persuade them to get an RSEP, which would slow things down quite a bit.

We also have the issue of scalability. So, for a relatively small ccTLD like .sk, it's easy for us to scan 400,000 domains once a day, but we are on the order of 20 million domains. Suddenly, the kind of infrastructure, potentially, we might need to support that is much higher.

As a back-end service provider, we want to make sure that our registry operator partners are happy with what we're planning on doing so we need to make sure we're communicating with them.

The real issue for us, I think, is in the gTLD world, registrars have an obligation to maintain a replica of the registration data that is up to date. Recent changes in registration data directory services, i.e. WHOIS and RDAP mean that, actually, the DS record does matter.

So, a registrar that just has to provide an RDAP service has to include the DS record information in the response while, on the Port 43 WHOIS, they don't. So, there is an additional obligation on the registrars to populate that data in their RDAP responses.

And so, if they want to be able to provide a reasonable RDAP service, they need a local copy of that and, we are making unilateral changes to the DS records, they are going to be out of sync.

There is a mechanism to notify them of changes like this, which is the Change Poll Extension. So, we would probably not look to deploy CDS scanning on our gTLDs until the Change Poll Extension is implemented, as well. So, we can tell registrars, "Hey, we just pulled a new DS record for your domain. You need to update your local record."

The other potential impact is that the number of name server operators that we would be talking to is much larger. So, the population for .dk is relatively small. It's quite a concentrated market. We're operating on 20 million-plus gTLD registries. Suddenly, we're talking to pretty much every authority service in the world. It seems to me like there is a high risk of operational issues and we want to understand what the potential risks would be if we started doing CDS scanning at a large scale on our gTLD platform.

STEVE CROCKER: Thank you. I'm going to take the [privilege to] ask you a question for clarification. You made reference to the Registrar Agreement, the RA, and about getting approval from ICANN. And so, I gather that you are operating both with the ICANN contracted parties and also, presumedly, with ccTLDs who are outside of that system. Can you say something about the mix or the totality? Just give a slightly clearer picture, here.

GAVIN BROWN: The vast majority of the domains now on our shared registry platform are gTLD domains. We do have a number of ccTLDs on that platform as well but we generally tend to treat them the same because the target audience for those domains is the same, as well.

In theory, we could turn the CDS scanning on just on our ccTLDs and worry about the gTLDs later but I think my preference would be to treat them all as the same.

STEVE CROCKER: Good. All right. Thank you. As I said, we'll go around and do a second round for comments, generally. Also, I see comments are showing up in the chat room. We'll come back to all of that but I'm going to move on, now, to Brian Dickson. Are you there, Brian?

BRIAN DICKSON: Yes, I am. Sorry, I was muted.

STEVE CROCKER: No problem.

BRIAN DICKSON: Hi. I am with GoDaddy and in the table we were listed as being a registrar but we're also a DNS operator so I'll have comments on both of those.

STEVE CROCKER: My apologies.

BRIAN DICKSON: That's okay. Part of that is that most or a lot of our DNS customers are registry customers, as well, but not all of them. So, we do have a substantial number of DNS customers for whom we are not the registrar. Okay. Next slide.

So, for our registry customers, we make the updates both via DS 3 EPP, as well as through CDS and CDNS key for all of our sign zones, regardless of what the TLD is. We've been doing DNSSEC for over nine years and all of our infrastructure is built around supporting all of that. For DS, it's always going to be the case that a customer can unilaterally update, if they're a registrar customer, the DS records through their UI, which gets pushed to the registry.

We generally handle DS creation and updates. Once DNSSEC has been enabled we do it as a managed implementation on the hosted DNS, so instead of DNSSEC-signed.

Making changes to the DS outside of that requires disabling the DNSSEC and then, for an end user, making manual changes. So, generally speaking, we're geared around automation. That includes both the CDS/CDNS key and the DS updates. Next slide.

Regardless of the TLD, we always do the CDS and CDNS key. From our perspective, that's the thing that we believe is going to scale the most, especially to enable and support third-party DNS operators. So, we really want to push this.

I think the only missing piece is the notification of changes to DS records from the registry to the registrar. I think there may be a distinction between the need for notification to non-CDS/CDNS key updates versus CDS/CDNS key updates.

So, if it is happening through CDS I'm not sure that there is any value. The reason is that DS records are, effectively, write-only. And so, a DS record is only used to populate the registry itself. So, other than the responses that are being asked for in RDAP, I don't think there is any value to actually keeping DS records themselves, except, potentially, as a way of validating what's in the registry. Even that, I think, is of dubious value.

That's my opinion. I'm not sure if I've given enough thought yet but I think that's something that we should collectively look at. If we decide that having those DS records in RDAP is actually of no use, maybe we can remove them from the RDAP requirement.

I think that may be a bit contentious but I think that helps solve most of the issues around whether polling and using the CDS and CNDS key is viable. Is there a next slide? Cross-signing, just a brief comment. We don't do anything yet but we're working on this. I think that's the end of my presentation.

STEVE CROCKER: Thank you. As I said, we'll come back around for some more comments and Kathy will read out the questions and comments that are piling up in the chatroom. Jothan.

JOTHAN FRAKES: Hi. So, I appreciate the privilege of speaking here today. I'll go through really quickly because I think … I operate a very small registrar but we do support DNSSEC. I participate heavily in a variety of the discussions so I thought I could offer some inputs here that may be a little bit different. Could you go to the next slide, Steve?

So, there is just some fundamental stuff about registrars and DNSSEC. The way the system is designed, registrars are the interface with the registrant and that was by design in the shared registry system. And then, key information gets furnished to the registry via the registrar, and that's how that design is, through the EPP.

You've got a separate registration and resolution chain and those work to provide the DNSSEC by working in harmony so that those records match. Now, the DNSSEC stuff is complex but there are a lot of benefits. We're starting to see some of the new innovations and technologies that really want to see DNSSEC in order to have a higher trust, such as authentication validation systems, and some blockchain things that are being developed also really require this as a platform. And so, it's a gateway into new and innovative things as well as a higher degree of trust.

So, we recognize the benefit and value that DNSSEC provides and we want to make sure that we can get those out to the customers. Pardon, I got muted. I think that there is [an inner] relationship between the level of adoption that there is with DNSSEC and its complexity and/or ability to get to market, but we see that that will probably be growing. Can you go to the next slide?

So, registrars have put in place some kinds of solutions. Typically, we have two paths in getting the information to the registry. There is either the model described by Gavin and others in RFC 8078 where the zone registry essentially pulls it, and that has some pros and cons. I'll let people enumerate through it.

It's not widely supported by registries and it does contain an overhead. The nice thing about it, though, is it allows a registrant to self-manage and the pull method bypasses any constraint on user interface of APIs at a registrar.

The automation makes it less prone to human error so that's all positive. The pace and frequency of updates, as Gavin indicated with a registry of a scale of 20 million, that might mean quite a few scans in a day and there is a lot to get in with respect to the frequency of updates. Next slide, please.

So, the other path, which is going through the registrar, which is either a manual or an API-based process, if that is available. There is a flavor of that, which is where the registrar has set up their own name server system in order to provide DNSSEC due to the complexity of those working together.

So, you've got the records being updated and synchronized for the registry and DNSSEC will work. Manual updates, though, are quite complex. We've got a variety, I think a spectrum, of users out there. Many are technically sophisticated; many also aren't, and they need the big, red, easy button to do this thing that automation can provide.

So, third-party providers may or may not offer good documentation on how to get the registrar to do this. Registrars are trying to make it work but the big challenge—and it's kind of highlighted there—is that registrant is going to want their support from the registrar even if there are issues that are spawned from a third-party provider.

The registrant will go to the registrar for support so we have to find ways to make that work well. Registrars may have to look at, are these third party services going to disrupt something that is already being offered to the customers? We want to make sure that that happens in a way that best serves the customer.

The automation method, there are N+1 integrations, so that means every registrar may have a different API that they offer. A lot of the APIs aren't closed to just manipulating the DNSSEC information, so there are some other areas that registrars are concerned about in getting third party access into this.

So, that kind of concludes this. I would just want to say that, as a third-party provider, many of the third-party providers see the complexity and are starting to accredit themselves in order to transfer domains to themselves, in order to offer the end-to-end solution.

And so, we start to get into slamming issues, as well, with respect to customer support. So, we want to make sure that, as registrars, we're doing the right thing and getting DNSSEC implemented, but there are some complexities involved in the system that we hope to get untangled through automation and some of the things that we'll be discussing.

As Brian mentioned, I see a lot of promise, potentially, in Domain Connect and some other things that we could use as standards to help us move forward. But of course, CDS or DNS key pulling, as defined in RFC 8078, does really help go around the system. We just need a way—

because we're compelled by the temp spec—to present, as registrars, the information. We're compelled to provide that so we need to be in-sync as a registrar with those actual records. That's essentially everything I wanted to share and I'll save for questions and answers to balance the time.

STEVE CROCKER: Super.

JOTHAN FRAKES: Thank you.

STEVE CROCKER: Thank you very much. Let's see. Oops. Ólafur, I need to do a quick switch, here. One second. You're on.

ÓLAFUR GUÐMUNDSSON: Hi, Steve. Hi, everybody. I'm Ólafur. I have been working on the DNS protocol but [inaudible] the DNSSEC part of it for a really long time. I've also been involved in TLD operations.

I am thinking that many, many, many years ago, we made a few mistakes. Undoing those mistakes is painful in the modern day, and I want to thank my panelists for actually highlighting some of the issues. Can I have the next slide, Steve?

So, when we started, at Cloudflare, to provide DNSSEC we were publishing information for all of our users on how to update registrars, to upload DS records, etc. We ran quickly into the issue that, basically, doing this was just like crushing the Berlin wall. Yep.

There is a privileged class of people that can do it very fast. Those are the ones who happen to be the registrars because they have access to the registration system and can put what [inaudible] anybody else. We have, like Jothan talks about, that registrars want to do the right thing but registrars also, in many cases, have what we call “resellers.” The consumers don't know whether they are talking to the reseller or the registrar.

And so, finding the right place, getting the keys accepted based on a new algorithm, was also [inaudible] at the time. We were promoting, for a while, the idea that the registrar should make available an API that third party operators could go and hit to have DS records uploaded.

Unfortunately, that went nowhere so we have thrown our whole weight behind the CDS and CDNS key automation. We would love to have the system to acknowledge that the NS and DS records are operational parameters and they should be classified differently than information like on, “Who is the registrant?” and, “who is the payment authority?” etc., and we would like to be able to adjust those in the case after DNSSEC is established on a domain.

The expectations from users and providers of information on the Internet is very different than Steve and all of those who are working on the [artifice] of the Internet. If things don't happen in a few minutes the system sucks. That is their attitude. So, getting things changed rapidly is the next thing we want to be able to do.

Automation will improve the accuracy and mistakes and avoid costly changes when things evolve. One of the things that we run into, like I mentioned before, is that when we started signing with DNSSEC we were one of the first to use a new algorithm.

It turned out that the hardest thing about letting information to registrars was that their user interfaces did not support or know about this algorithm. That caused lots of angry calls from our users to registrars where people we talking past each other because they didn't understand that the algorithms were different. We apologize for that. Next slide, please.

So, if we think about this in the DNS ecosystem for the 21st century, if we leave behind the DNS of the 20th century, we have [become the] registrar also and, for our domains, we push the DNS record into the registry, so that works perfectly fine.

What we do when our users turn off a DNSSEC or decides to leave us, we remove the DS immediately. We are also deleting any DSes there are when the domains are transferred into us because we have noticed that a number of transfers happen and fail or cause operational problems because there is a DS that has the old keys in them. That is very annoying and needs to be fixed, possibly by an ICANN mandate that that DSes is always deleted on transfers. I know Steve disagrees with me on that one.

From the DNS authoritative side, we sign on the fly so that we can change information really rapidly. We publish the CPS and CPNS key for all of our zones and we are 100% in support of the [inaudible] which tells people to delete the DS record. When users disable the DNSSEC we keep signing the zone until we have noticed that the DS record is gone.

So, full automation/full-scale scanning is, I think, the only way forward in a system that is built on the side, whether it is Domain Connect, API calls can only help deal with real-time issues when the DNS providers and the registrar know about each other. But when you are dealing with, let's say, registrars in Slovakia, yeah, we may not know about them even though the Slovakian registry knows well about them.

So, automation, automation, and reclassification of information that is stored in the registration system of what is operational and what is account-related so the operational information can be changed to be an automation. That's it. Thank you.

STEVE CROCKER: Thank you very much. As I said, we'll go around quickly to each of the panelists. Kathy, how much time do we actually have, here?

KATHY SCHNITT: We have just shy of 20 more minutes.

STEVE CROCKER: Oh, cool. So, we have time for a Q&A from the audience, as well. All right. Again, I'll just go in order.

BRIAN DICKSON: Sorry, do you mind if I jump in very quickly?

STEVE CROCKER: Go ahead.

BRIAN DICKSON: One thing I forgot to mention and has come up in some of the subsequent presentations was the issue of potentially using Domain Connect. What was originally suggested by one of my colleagues was Domain Connect for DS updates. I want to clarify we're only suggesting an extension to Domain Connect for the initial DS, not for maintaining DS. That's really only more like a trial balloon to discuss. If it doesn't happen it wouldn't really be terrible because it could be difficult to implement but at least it's a starting place to talk about how to get the original DS in.

STEVE CROCKER: Thank you. I appreciate it. Okay, good. Well, you've just used up your slot. I'll just come back to the beginning. Jim, do you want to add anything?

JAMES GALVIN: Yeah. I think I just want to comment on the issue of DS records versus key records. There has been some discussion in the chat on this issue of DS records and whether or not it is indicated in RDAP output, for example. It's also part of the Q&A that comes up later.

I think that it's an interesting question to be asked in this provisioning question, how much control you want to allow and to whom. What does the DNS provider get to control and decide versus what you want the registry to control and decide or the registrar to control and decide?

I think those are important points of control that have to be determined/ have to be agreed to. Maybe there is more than one solution that could be accorded, here. We're working in a limited ecosystem at the moment but I think that's the operative question that we have to ask ourselves.

The DS records are published in RDAP and WHOIS because it's a nice [outer ban] mechanism for being able to validate to see that information. Gavin talked about the idea that the registry doing the polling and reaching out. Well, that can potentially be an issue when you start talking about this at scale. And so, you really have a different set of problems to think about when you have that point of control that changes.

And then, there is just the general authority question in our ecosystem about who really owns the data – the registrar, or the registry, or the DNS provider? That's in important consideration here. Thanks.

STEVE CROCKER: Good. Irwin, do you want to say anything more?

IRWIN LANSING: I've actually got two points. Let's talk about one. To me, one of the big issues here is, in this case as the registry, how do we know who actually is the name server operator who gets to register the name server? It was on my first slide. If it's a name server inside of a .dk domain, we can go to a registrant and say, "Is this your name server?" but for any other TLD we don't know who is actually behind that name so anyone can, basically, go in and register any domain name and add a TLD. I can go out and register Cassandra.cloudflare.com in the "dk" registry.

Now, I can't really use that yet, at that point, because the privileges come from the domain names that are delegated to that name server. So, before any domains are delegated there is no use for it.

But I can see some scenarios where that can be exploited. So, if we assign that DNS operator some data they can manipulate at the registry level, how do we make sure that we’re talking to the right DNS operator?

STEVE CROCKER: Interesting. I'm just going to keep going. Gavin, do you want to add anything?

GAVIN BROWN: No, I don’t have much more to add. There are some questions that have been put in that have been directed at me. Maybe we can do those later? I don't know. No, nothing more to add.

STEVE CROCKER: Okay. Jothan? Are you still there, Jothan?

JOTHAN FRAKES: Lost audio there for a moment. Yeah. I think one of the interesting things we’re seeing is that third-party DNS providers and ccTLDs not being subject to the same regulatory constraints governing what they have to do have the flexibility to do some creative things, here.

When you get into some of the obligations that contracted parties have with respect to RDAP and even the life cycle and disruption that are compelled by our agreements, the requirement that we offer DNSSEC in the 2015 registrar agreements, there is a variety of things that we are compelled and must do. A third-party DNS provider is not necessarily under that same constraint.

And so, if we look at them in silos it probably seems quite simple, but when we start to look at the holistic system we have to think about the things that are inter-related, that if you change this dial you'll have to change this dial over here.

So, if we did say that things did get pulled using DS—just spit-balling a particular concept—that would mean that we have to go and look at RDAP and the requirements on registrars under the temp spec and/or the future spec, that registrars have to show that information versus the registries showing that information.

If we were to take and look at a third-party provider, I think you'll see little resistance from registrars or light resistance, other than those RDAP obligations, with respect to DS for key records by registrars if they're updated outside that system.

Where you'll hear a lot of resistance is the updates to the DNS servers themselves for a given domain name. You'll find a little bit more resistance, and that's largely in respect to the obligations that are there governing and enforcing compliance.

So, there are a variety of things that have to happen. For registries, they've got our steps, as Gavin mentioned, and some scale concerns if we did go the DS. So, these were complex. Now, they're complex on their own but they also have to be considered in the larger scheme of the regulatory web that we find ourselves caught in. That's my closing comment.

STEVE CROCKER: Thank you. Ólafur?

ÓLAFUR GUÐMUNDSSON: Nothing to add.

STEVE CROCKER: Okay. We have, as I said, quite a few comments and questions in the chatroom. Kathy, can I turn this over to you? But before you begin, what I'd like to do as a result of this session is capture it all and pursue it on a mailing list. Let's see. Mailing lists, he said. What mailing list are we talking about? Whoops. There we go.

So, there is a mailing list at dnssec-provisioning@shinkuro.com, which has got a couple dozen people discussing, basically, these issues, and quite focused. Anyone is welcome to join. I think you have to send a message to me. I haven't quite mastered how to make this as automated as it should be. I'm happy to add you.

It has got no status, no authority—it's really an IETF-like design team to think through the issues—but I hope that the output is a sensible conversation on this topic and some focused attention on possible ways forward. With that, Kathy, let me turn things over to you for what we have time for from the chatroom and from anybody else who wants to ask a question.

KATHY SCHNITT: Thank you, Steve. Our first question is from Francisco; question at Gavin. “Do you allow EPP updates to DS/DNS key records in .sk? If so, what do you do if there are discrepancies between the EPP update and what you see in the CDS scanning?”

GAVIN BROWN: So, CDS scanning doesn't preclude changes to DS records via EPP and a change via EPP will overwrite the DS records in the registry if they differ from what was there previously. However, if the CDS record differs and cannot be validated then the data provider by EPP will prevail because if we can’t validate the CDS records using the DS records provided by EPP then we will ignore them.

KATHY SCHNITT: Thank you, Gavin.

STEVE CROCKER: Gavin, a natural question as I listen to all this is, what happens if somebody updates the DS record directly through EPP but the records persist in the zone, the CDS records? Won’t they be picked up on a later scan and then override the direct update?

GAVIN BROWN: So, there could be a scenario where there is an overlap. Obviously, during a rollover or a double-signing process, it may be possible that we can validate the CDS records using the DS record we have. If that’s the case, that we are able to validate it, then the CDS records would get reapplied. But if we can't validate the records then what is provided by EPP prevails.

[ÓLAFUR GUÐMUNDSSON:] Yeah. Steve, this is one of the automation problems. If the information that is somehow incorrect gets into the system, unblocking it is going to require a human action unless we have a policy of what to do.

GAVIN BROWN: Yeah, thank you. Back to you, Kathy.

KATHY SCHNITT: Thank you. Next question for Irwin and Steve: "By requiring name servers to be pre-registered with the registry, does .dk require the registrants to operate/use/host their domains only within the registered NS space?"

IRWIN LANSING: Yeah. So, the name server has to be registered within the registry before any domain can be delegated to it. So, for the registrant—and this is another one of the caveats with this model—there is a scenario where the registrant goes to a registrar to register a new domain. The name servers are not there yet, so when the registrar sends the application to us we can only refuse the application.

Then, the registrar has to go back to the registrant, to ask the registrant to go to the name server operator, to register the name server, and then come back to the registrar to start all over.

STEVE CROCKER: I'm not sure whether this was intended as part of the question but do those name servers have to be within the .dk bailiwick or can they be anywhere on the net?

IRWIN LANSING: They can be anywhere. Within the "dk" bailiwick we ask the registrant to approve if it's not the registrar itself registering a name server.

STEVE CROCKER: Thank you. Kathy.

KATHY SCHNITT: Thank you. The next question is from Peter and there is a plus-one for Peter's question from Marcus. "For 'sk,' the slide showed a correlation between CDS scanning and DNSSEC growth. Is scanning the only way of provisioning? If not, what is the CDS versus EPP ratio?"

GAVIN BROWN: Okay. So, as I mentioned previously, EPP can be used to add DS records to domains/.sk. The question about what the ratio is, the slide shows you the ratio. We were hovering at around about 3% or 4%

adoption in .sk before we turned CDS scanning on, and then, a few months later, we're at 35%, so an approximate and order of magnitude difference in terms of DNSSEC adoption.

I haven't gone through the exercise of working out why. If your assumption is that most domain names use the DNS of the registrar, and if that many registrars are doing DNSSEC for their domains, why weren't they adding the DS records before we turned CDS scanning on?

One possible answer to that question would be it was simpler just to wait. They knew that CDS scanning was coming and then, therefore, they said, "We won't bother implementing RFC 5910 because we know that CDS scanning is coming and, therefore, we can just publish the CDS records and be done with it and don't have to worry about implementing that APP extension and the internal communication channels from the DNSSEC signing system to the registry interface module."

So, it may be an interesting exercise to work out why all of those zones that were secure weren't secured before we turned CDS scanning on. That might be an interesting insight to have on some of the challenges that registrars and DNS operators have.

STEVE CROCKER: Good.

KATHY SCHNITT: Thank you. Our next question is from Duane: “For .sk, do the registrars do CDS scanning or is the registry scanning?”

GAVIN BROWN: So, the registry does CDS scanning. That doesn't preclude registrars from doing it themselves, as well. It would be straightforward to write something that does a CDS query and then turns that into an EPP command to update the domain. That was the deliberate design decision on the part of the RFC that defined CDS and CDNS key records.

KATHY SCHNITT: Thank you. Our next question is from Jacques: “Where is it defined RDAP for gTLD needs a DS record, and why?”

GAVIN BROWN: I'm happy to answer that question. So, the RDAP operational profile says that the registry or registrar must populate those fields. So, I'm just going to read from section 2.10 of the response profile: “The domain object in the RDAP response must contain a secure DNS member including at least a delegation signed element. Other elements, e.g. DS data of the secure DNS member, must be included is the domain name is signed and the elements are stored in the registry or registrar database, as the case may be.”

JAMES GALVIN: I will add to that. I gave a little bit of history in the chat room. I have answered this question, also, and I said it in one of my earlier comments, too. There are two important distinctions to make, here. One is that the fact that it's in the RDAP response does not mean it has to be displayed. That's an important thing to keep in mind, okay? It is required to be returned by a registry or registrar but it's a client responsibility to figure out what to do with that. So, if you don't like it just get yourself a client that doesn't display it.

For legacy reasons, I'll point out that it was always included in WHOIS and originally, in the early days, it was put out there in WHOIS because it provided a nice out-of-band way in which to confirm that your DS information was actually correct in its transfer from the registrar, to the registry, and then to the DNS.

You could compare your DNS from your provider with what the registry has in its DNS system. Whether or not this is valuable is, I suppose, kind of an open question, which is why I made my earlier comment about whether or not you display it is your own decision, because that's a client thing. But we're carrying forward this legacy position of being able to see that in an out-of-band way for those who are interested in it. Thanks.

KATHY SCHNITT: Thank you. We have time for one more question and after that any comments and questions that were not answered I will go ahead and send to the panelists and we can reply to you via e-mail.

STEVE CROCKER: I have a question before you go. Kathy, I have a question. This session is recorded. Will there be a transcript of it?

KATHY SCHNITT: Yes, correct. There will be a transcript and it will be posted on the public site.

STEVE CROCKER: And these questions in the chatroom will be retained, as well?

KATHY SCHNITT: Correct, yes.

STEVE CROCKER: Super, thank you. Go ahead.

KATHY SCHNITT: You're welcome. Our last question comes from Mark: "In EPP, doesn't the registrar notify the registry that they are DNSSEC-capable? Thus, not all domains would need to be scanned for now."

GAVIN BROWN: Sorry. Do you want a response?

JAMES GALVIN: Sure. I'll go first this time. I answered this question in the chat, also, and just commented that there is no flag, per se, between the registry and the registrar. So, a registrar doesn't announce the fact to the registry that it has that information. So, it's a local registry policy, in fact, whether or not you have to do anything special between the registry and registrar to provide that information.

And I just offered that we as a registry, Afilias, do require our registrars to go through an OT&E where they give us this information, just to confirm as we do with everything when we accredit a new registrar. But once that's there, we just take whatever they give us.

And then, it's the presence of the DS information in the registry database that we would use as a flag if we were to want to make decisions based on that. But there is, otherwise, no real out-of-band flag that is a part of that. Thanks.

KATHY SCHNITT: Thank you. Oh, go ahead.

RUSS MUNDY: Yeah. I have our closing and we have roughly two or three minutes. I see that Kathy has the slide up on the screen and we do truly, truly want to get feedback from people. In some ways, I almost feel like we want to have people think about a LinkedIn or a Facebook kind of post. The more reviews we get, the better.

But seriously, folks, please do give us feedback. We were going to try to have a few questions but we've had such a lively discussion we won't have time for any more. So, please go ahead and post them in the chatroom and, if you have other preferences besides the review session, you can probably make use of our paper submission or idea submission "call for participation" mail list, which I've also just put in the chatroom. So, if you want to send further feedback there then please go ahead and use that list, also.

I wanted to express my great appreciation for all of the participants today. We were well over 100 at one point; all the panelists and all of the people who helped organize the panel. And from someone who has been involved in this for a long time, I am very pleased with the results of our first virtual DNSSEC and Security Workshop. I think it went very well.

But the biggest thanks of all goes to the ICANN staff, Kathy in particular, but there were also others helping a great deal to make sure this went as smoothly as possible. With that, thank you, everybody. Kathy, was there anything at the end you need to [cover?]

KATHY SCHNITT: Yes. Thank you, Russ. First of all, I want to thank all of the panelists who joined today, especially at some crazy hours for some of you. We appreciate it. I want to thank the DNSSEC Program Planning Committee. They worked very tirelessly and very quickly to make this happen virtually.

My tech support guys have been fabulous behind the scenes to make this run smoothly, and my colleague, Kim, for being my back-up for this entire session. Without her, it couldn't have been possible. Other than that, this session has ended. Thank you.

DAN YORK: Thank you, everyone, especially those of you who stayed in the whole time. Thank you.

RUSS MUNDY: Thanks, everybody. Bye, now.

[END OF TRANSCRIPTION]