ICANN69 | Virtual Annual General – DNSSEC and Security Workshop (3 of 3)
Wednesday, October 21, 2020 – 16:30 to 17:30 CEST

UNIDENTIFIED MALE: Good afternoon, good evening, good morning, everybody. Our last session for today is all being taken by one person. That's Willem. Willem is the developer and researcher at NLnet Labs where he works on open standards and opensource software for core Internet protocols. He's especially interested in delivering first class security and privacy with DNSSEC and DNS over TLS to end users at the edges of the Internet. He is passionate about his work and cannot help himself talking, explaining, and presenting about it. So over to you, Willem.

WILLEM TOOROP: Thank you. Let's see. Share my screen. There. This presentation is about a hackathon project that I did with a bunch of other people in the spring of 2017, but the project is still running and growing and has played a role in several different research projects. A few weeks ago the project was mentioned a few times in the email thread on the DNS operations mailing list of DNS-OARC which was noticed by someone from the program committee of this workshop. I had already submitted another presentation for this workshop, but they asked if I could present on this too. I think that's actually a good idea because the research in the second presentation is also assimilated into this hackathon project.

In April 2017 I did this hackathon project not on my own. It was a team effort. The goal of the project was to provide insight into caching resolver capabilities that are in use on the Internet. Our approach was to use RIPE Atlas to schedule all kinds of different DNS measurements and then process the results.

For the ones that don't know what RIPE Atlas is yet, it's a measurement network consisting of 11,000 probes. Those probes were all handed out by so-called RIPE Atlas ambassadors at conferences and other events with the request to take them home or to your office and hook them up to your network. If you have the probe, you get RIPE Atlas credit which you can then use to run measurements yourself using other probes in the RIPE Atlas network as vantage points.

The capabilities that we wanted to measure were things like can the resolver reach that the authoritative of IPv6? Can it return over TCP if requested by the authoritative [inaudible]? Does the resolver do a DNSSEC validation? And if so, for which algorithms? The QNAME minimization, EDNS Client Subnet, everything you ever wanted to know about caching resolvers but never dared to ask.

With the results you get from RIPE Atlas, they are from the vantage point of the RIPE Atlas probe. The default viewpoint is like you're a user in the network which hosts that specific probe. The probes also get DNS resolvers from [the HTTP] just like other devices in that network. The results you get from RIPE Atlas are those that you see from that vantage point. For some measurements this is fine, but for

others it's helpful to have the vantage point from the authoritative side as well or both.

For those doing the hackathon we created a special authoritative name server. Actually, Jerry Lundstrom from DNS-OARC did that. The authoritative name server, the DNSThought Daemon.

Here you see the DNSThought Daemon in action. We are requesting a query that does a measurement whether the resolver supports TCP. We're targeting Quad9 and it replies, so it was able to get to the DNSThought Daemon of TCP.

Besides getting the answer signaling that this resolver indeed has TCP support, DNSThought also puts the IP address of the resolver itself at the authoritative side in the answer. So this IPv6 address is actually an IPv6 address in use by Quad9 resolvers.

Now we have these perspectives. We can see the IP address is configured as resolver which it learned from [the HTTP] which could be a local resolver or a public resolver. Also, it knows the IP address as seen at the authoritative site through the answers that are returned by DNSThought Daemon.

With these vantage points we classified three types of resolvers. If the autonomous system numbers for the IPs of the probes, the resolver, and the authoritative are the same, then we classify such a resolver as internal. If the autonomous system number for the IP address of the resolver is different from that one of the probe but the authoritative IP address is seen with the same autonomous system, then we call this

an external resolver. And other combinations for which, for example, the IP of the resolver and the IP at the authoritative and the IP of the probe are all different agents, we call that forwarding.

We didn't just schedule tests for static zones with things that DNSThought Daemon returns. We also incorporated existing tests, such as this test which came from internet.nl showing you if your resolver has QNAME minimization support.

These were the measurements that we scheduled during the hackathon in 2017. They were scheduled with special superpowers from Emile Aben that made sure that they targeted all probes in the RIPE Atlas network and that measurements will continue to go on forever and ever and also will incorporate new probes that come into the RIPE Atlas network. At the time we also created a portal that showed you the status of the resolvers of the different probes.

And so the hackathon project was finished, and then in June I started participation with Roland Van Rijswijk-Deij on the Root Canary Project which has a goal to mentor the upcoming DNSSEC key signing key rollover.

For the Root Canary we set up an infrastructure to monitor validation support for all algorithms more or less as a side project. For this we created a matrix of zones signed with all DNSKEY algorithms times all the different delegation signer hashing algorithms.

The zones all have both secure and on purpose bogus address resource records. They are all subdomains of through rootcanary.net

domain. The zones use a name scheme which refers to the DS algorithm used to delegate the zone, the DNSKEY algorithm with which the zone is signed, and whether the zone uses NSEC or [NSEC3].

We scheduled measurements for all the secure and bogus records for all those zones again with Emile's superpowers. The results graphs of DNSThought measurements are [housed] on the website dnsthought.nlnetlabs.nl. That also contains a reference to all the RIPE Atlas measurements that we used with the project. They are all public measurements, and others can use those results from those measurements directly too.

Root Canary also had and still has an online algorithm test using these zones which you can still do showing you what DNSSEC algorithms the resolver that you are using supports. This can be found on rootcanary.org.

And then almost a year later in July 2018 Moritz Miller, who presented earlier in this workshop, joined too to monitor the root trust anchor rollover. We scheduled new measurements for the key signing key sentinel queries. It's a special feature of validating resolvers in which they reveal which trust anchors they have with which root trust anchors they can validate DNSSEC names or DNSSEC validated domain names.

Then in October that year the actual KSK rollover happened, and we scheduled another two measurements monitoring the root's DNSKEY seen at all the different resolvers. We made a graph for this showing how the new key was distributed over the RIPE Atlas network or how

the resolvers on the RIPE Atlas, how quickly they learned about the new key so to say.

All this collaboration resulted eventually in a scientific paper presented at the Internet Measurement Conference in 2019 in Amsterdam. At that conference we received a distinguished paper award for it.

For the rest of this presentation I will show you [inaudible] of different plots that can be seen on the DNSThought website. Note that above the plots that I am showing there is a link linking to the DNSThought website where you can see these specific plots.

This plot shows the distribution of internal, forwarding, and external resolvers in October 2018. Just after the root KSK rollover there was a DNS-OARC workshop in Amsterdam, and this is what the distribution looked like back then. I'm showing this because the distribution is quite different now. It's like this. So the group of external and forwarding resolvers on RIPE Atlas increased from 22% back then to around 30% now at the expense of internal resolvers, which is interesting.

This is a plot showing the top ten autonomous system numbers of the IP addresses of resolvers as seen from the authoritative sites. You can see that Google resolver is the most popular resolver on RIPE Atlas, followed by CloudFlare, followed by OpenDNS, etc.

On the DNSThought web pages it's also possible to select a certain property or capability and then view all the different plots only for the

resolvers which have that property or capability. In this plot we have selected that the resolver should be internal, have that property. And we are also looking at a plot showing internal, forward, and external resolvers. So we see 100% internal resolvers as expected. We can clearly see the decline in internal resolvers on RIPE Atlas.

This is the top ten ASNs seen at the authoritative but with internal resolvers selected. I think what's interesting here is to see that the ASNs for internal resolvers are almost all ISPs: Comcast, Deutsche Telekom, Liberty Global, etc. Whilst as you look to the forwarding and external resolvers, they are all cloud DNS resolvers.

This is the plot at this moment for the forwarding resolvers. It's nice you can also clearly see how new cloud DNS resolvers are introduced and how they are taken up by the network, for example, the Quad9 that started on 16 November 2017 and the Quad1 on 1 April 2018.

These are the ASNs for external resolvers, the top ten ASNs. You can clearly see that there are less and less different resolvers. Overall we can see that the number of ASNs seen at the authoritative decreases if you go from internal, forwarding, and to external.

This plot shows RSA-SHA256 support of the resolvers on RIPE Atlas right now. Currently this is 64.2% which is amazingly high. I heard Dan York say this morning that worldwide it's 25%, so this tells you something about bias in RIPE Atlas. Of course it's mostly European based [approach] and also quite a few in North America but not in the rest of the world. I also heard them say that in Europe DNSSEC

validation is now 43%. So RIPE Atlas clearly still has a better capability resolver bias.

Also note that RSA-SHA256 is the algorithm used at the root. So this is the base algorithm you need to be able to validate to be able to DNSSEC at all. So it's interesting to have a look at all the other algorithms if you look at resolvers which support this algorithm.

It's actually even interesting to see that RSA-SHA256 support when you're only looking at resolvers that have this supported because next to the big pie chart which shows that all the selected resolvers are indeed supporting RSA-SHA256 there are three small ones that show you percentage of probes that have a resolver with this capability and also the other values for this property.

So all those probes, there are 7,500 probes that have at least one resolver that can validate DNSSEC; 10.5% of those also have a resolver configured that does not do validation. So in the end that means that only 60.1% is protected by DNSSEC validation because if such a probe cannot get an answer from the DNSSEC validating resolver, it will fall back to the non-DNSSEC validating. So that's also interesting results you can see on the DNSThought pages.

Here are a few plots for other DNSSEC algorithms for the resolvers that support DNSSEC RSA-SHA256. In the upper left corner are two pie plots for ECDSA-P256 and ECDSA-P384 which is almost the same as RSA-SHA256. Below that we see support for DSA which is now [inaudible] to not support on the resolver. So we can indeed see that less and less resolvers support this algorithm. On the right we see the

uptake of ED25519 and also of ED448. These measurements have also been used in the DNSSEC algorithm agility study that Moritz presented earlier today.

Another interesting graph that I wanted to show you to show how quickly a new feature is taken up is the ASNs that support the root key trust anchor sentinel mechanism in which resolvers reveal which root trust anchors they have. We are now at 20% of resolvers at RIPE Atlas support this algorithm. You an also see that even after the KSK roll for which this mechanism was actually developed there is uptake of this mechanism recently with Verisign and even more recently with Free SAS, which is a French ISP.

We can also see which resolvers still support the old DNSSEC trust anchor which is interesting, I think. There are still a few, but less than half a percent of the resolvers [inaudible].

With DNSThought you can also discover strange things. When looking at the DNSKEY algorithm the top agent for support of RSA-SHA256 I noticed this dent. The dent was caused by CloudFlare falling out of the list of top ten ASNs. Selecting all the results of the DNSSEC the properties for resolvers AS13335, the autonomous system number of CloudFlare, confirmed that they indeed had a period in August 2018 in which at least Root Canary zones were not validated.

On that same page lower where it shows the support for the delegation signer hashing algorithms I noticed this red bump that indicates that some probes have broken delegation signer GOST support. Selecting the page that shows broken GOST hashes or GOST

delegation signer hash algorithm support and looking at the top ten ASNs we [failed] Google. They, just like CloudFlare, also confirms the incident. Google actually uses DNSThought to monitor the page there of Quad8 on the Internet and to monitor how new instances of the software is deployed [on the Anycast nodes].

At that time the ED25519 algorithm was only supported by CloudFlare and Google or almost only supported by those two. Looking at that graph [selecting] for this property enabled me to show the two incidents side-by-side.

We have other properties that we do with DNSThought too, such as monitor for EDNS Client Subnet, which resolvers do that and also which masks they expose to the authoritative server. Another interesting one is QNAME minimization. QNAME minimization, as you can see, is still growing, support for it. Very recently also Deutsche Telekom, very good. QNAME minimization measurements were also used in a scientific paper, and that paper won the Best Dataset Award at the Passive and Active Measurement Conference held in March 2019 in Chile. A nice side effect of this hackathon project is that it puts my name on all those scientific papers which win these prizes.

That's my presentation about DNSThought. Shall I just continue with the second presentation and leave questions for after that?

UNIDENTIFIED MALE: I think so. I don't see any questions or comments so far. You've got all of us by the throat, so I think just carry straight on. That gives you [33] minutes. So next presentation.

WILLEM TOOROP: Yes, indeed. I'm out of that [inaudible]. I'm a little short on time, so start presentation. Now I have to share my screen. There. The Current State of DNS Resolvers and RPKI Protection.

This presentation is about a research project by two students from the System and Network Engineering Master from the University of Amsterdam, Erik and Marius, who looked into route origin validation. How route origin validation is protecting DNS resolvers on the Internet. They did this research at NLnet Labs in January 2020. The goal for me with this presentation is to stimulate you to think about RPKI and the role it has in the dynamics of DNS.

So what's this about? Well, we DNS people have provided security means, DNSSEC, by which you can be sure you have the correct address for [certain name] provided the domain name is DNSSEC signed and the requestor is DNSSEC validating, of course. But DNSSEC does not help when connecting to this address. The looked up IP address itself can be trivially hijacked.

How? Well, suppose user Nelly from Network N would like to connect to her [house] Goofy in Network G, and Goofy has autonomous system number AS64503 with which it announces its prefix 8.8.8.0/23. The other [inaudible] learn this prefix by way of a path vector protocol

which is called BGP. In this way Network N knows that for this destination it has to route through AS64500 because this is the shortest path to Network Goofy.

But a malicious AS may announce a prefix it does not own. Traffic will be routed to the malicious AS if the path is shorter or if the malicious AS announces a more specific route. So in the picture Nelly's traffic for Goofy will be hijacked by the Malfoy Network M with AS666 because AS666 announces a more specific prefix from Goofy's IP range. It announces a /24 and not a /23. Also, the path to Malfoy's network is shorter than to Goofy. So your traffic to host Goofy will end up with Malfoy.

The Resource Public Key Infrastructure (RPKI) tries to deal with this by letting network operators cryptographically sign and validate prefix and origin data with so-called route origination authorizations (ROA).

How does that work? Here on this slide Goofy has created a ROA which states that prefix 8.8.8.0 with prefix length 23 announced from AS64503 is valid. This statement is cryptographically signed with Goofy's private key and can then be validated by Nelly's network and with the help of Goofy's public key. This is called, by the way, this validation process route origin validation (ROV).

RPKI does not protect the complete BGP path. AS64501 and AS64502 also need to do route origin validation for it to work. AS64504 might still route Goofy's IP address to Malfoy's network because it does not do route origin validation and it is a very short path to a more specific.

There's another [inaudible] which does protect the complete BGP path which is called BGPsec, but this [inaudible] is less popular.

So perfect. We have DNSSEC protecting the name to the IP address mapping and RPKI making it more likely that we connect to the right address. But what does this have to do with DNS resolvers? Well, Nelly's resolver does not to [inaudible] in there by RPKI protected Network N. She might use a resolver from AS15169 which does not do route origin validation. So even though Nelly is protected her resolver is not and Malfoy might hijack the authoritative name server of Annie's Network A and return a wrong IP address for Nelly's DNS lookup.

So even though Nelly is protected she might connect fraudulently to Malfoy's network on a resource owned by Malfoy because Malfoy managed to hijack the authoritative server for a DNS zone containing the [house] Nelly wanted to connect to. So DNSSEC can prevent this, but both the zone has to be signed and also Nelly needs a DNSSEC validating resolver. So a nice academic [attack] that I have sketched here, but does this really happen?

It actually does. A well-known example of this kind of hijack is the Amazon Route 53 hijack on 24 April 2018 which took over the myetherwallet.com domain with the goal to steal Ethereum cryptocurrencies.

You might already have noticed that RPKI and DNSSEC are fundamentally different beasts. For both you need the resources signed and the one contacting the resource validating. But with DNSSEC that's enough. It doesn't matter how you got it. If the resource

is signed and you are validating, you can be certain it is authentic and complete.

RPKI on the other hand can still suffer hijacking if another AS on the path is not validating. Also, it does not protect the [carried] traffic in any way so it has no integrity protection of the content of the traffic and also no authentication. The only thing it actually does is increase assurance of delivery to the intended AS. What you can say about RPKI is that the more of it there is in the world, the better it will work. Also, shorter BGP paths will help.

At NLnet Labs we have the University of Amsterdam across the street and if I am individual in a certain topic, I can invite students from there to look into that. They benefit from it because of the research experience, and I benefit from it because the work to look into a topic actually gets done and increases the probability of getting answers.

For this I invited students from the System and Network Engineering Master of the UvA. Erik and Marius took the challenge and they formulated the main research question as follows: What is the state of Route Origin Validation (ROV) on DNS resolvers? As a separate research question they formulated: Does the length of the AS path matter? How does Anycast influence the protection?

To test they set up two authoritative name server for two domains. One of the domains, valid4.rootcanary.net, is served on a valid prefix and the other domain, invalid4.rootcanary.net, is served on an invalid prefix. The invalid prefix is invalid because it has a ROA associated that

says that the prefix is announced by AS0. This AS0 does not exist, and also it does not announce this prefix.

The valid domain has a DNAME pointing to the invalid domain. The invalid domain has a wildcard A record. In this way a query for a [certain] random name which is [inaudible] of the DNAME in the valid domain would be ultimately answered with that address on the wildcard, but only if the resolver is not protected by route origin validation. Only then does it enter at the invalid authoritative name server. Furthermore, the random name below the DNAME will be seen at both the valid and the invalid authoritative.

So a RIPE Atlas measurement was scheduled with all probes sending to all the [inaudible] with a random ID. So here you have a closer look at what the measurement looks like. The DNS specific settings here show the query name and type and also $r-$t-$p which will be replaced with random hack string [inaudible] probe ID.

So this is how the resolution works in detail. If the identifier [random is timestamped] [inaudible] is seen at the valid authoritative but not on the invalid, then the resolver is protected by route origin validation. If it's seen by both the valid and invalid, then the resolver is not protected. Furthermore, those authoritatives are housed [on a beacon] that does its own BGP. So we also know the AS path [toward] those authoritative servers.

In [practice] we see this because resolvers don't like if an answer does not get answered. So it reveals a little bit about the internal resolver

infrastructure of the resolvers used at RIPE Atlas which is very interesting.

The beacon with the valid and invalid prefixes was provided by Job Snijders from the [entity]. The RIPE Atlas measurements again scheduled with the superpowers from Emile Aben. So thank you, Job, and thank you, Emile.

Erik and Marius did these measurements reliably from 23 January until 3 February. Timing was really good because in those 12 days the number of probe resolver pairs that did route origin validation rose from 7% to 15%. The right graph shows how well RIPE Atlas probes were protected. Probes have an average of two resolvers configured, and on 3 February 17% of those probes had at least one resolver which was route origin validating. However, one-third of those probes also had a resolver configured that did not do route origin validation. So the fully protected probes on 3 February was 11.5% of the population on RIPE Atlas.

On the left the ASes of the top ten most popular resolvers on RIPE Atlas are shown; 30% came from AS15169 which is Google. [But] all those queries reached the authoritative on the invalid prefix. The second most popular AS is 13335 of which only one-quarter reached the invalid [auth]. All the other popular resources did not do route origin validation except for AS12322, Free SAS, which is a French ISP.

On the right we see the ASes responsible for protecting the most queries. So CloudFlare number one, [inaudible] is one-quarter, and all

the others have full protection with Free SAS as number two and AS3265 which is Xs4all as number three.

So to answer the second sub research question: Does the length of the AS path matter? The ratio of protected queries from [the various received] BGP path lengths were counted, and this graph shows the number of protected and unprotected queries. The graph suggests that there is no clear correlation between path length and amount of protectiveness.

The second sub research question was: How does Anycast influence protection? For this the distinct CloudFlare prefixes were counted. [Assume] each Anycast node uses its own prefix. There were 160 CloudFlare prefixes counted in total. During the measurement period we could clearly see that CloudFlare was deploying route origin validation on the various Anycast nodes. And on 3 February half of those prefixes were protected.

These measurements are now also integrated into DNSThought which we can see how protection of route origin validation grew over the last ten months actually. After February it did not grow a lot. It more or less remains stable. You can see that CloudFlare is still the main influencer of route origin validating resources but also Comcast has recently begun. What's also interesting to see here is what Free SAS is doing.

Oh, and I also created measurement for IPv6, of course. These are the results for IPv6. There's a lot less protection of RPKI protected resolvers to IPv6 destinations. This is the AS12322. If we select in

DNSThought for this AS, you can indeed see that it sometimes does route origin validation and sometimes does not.

Here's CloudFlare. On 3 February it was about half that was protected and this is still the case. In January Emile wrote an email to CloudFlare to ask about this, why not everything was protected. And he received this answer in which they state that they also keep a default route at the Anycast nodes. So the appear locally, but for transit they might use a default route which might or might not be protected. They suggested to run a TXT bind.hostname query to all the Quad1 resolvers to see because it reveals which Anycast PoP of CloudFlare it targets. And then [inaudible] correction that it had to be the [ID.server TXT CH] query. Indeed, if I run this query, you can see I reach the PoP in Amsterdam.

These are all the PoPs. Green means is protected route origin validation. Red perhaps also but not for our beacon which is located in the United States, by the way. This is what it meant for the probes in the different countries. We used the same graph for the PoPs for IPv6. And here the probes in the different countries, how they were protected.

So future improvements. We looked at authoritatives only. Measurement network with more vantage points is definitely possible because it doesn't require the vantage point of the user per se. I would like to have more beacons all over the world to further investigate this thing with Quad1, for example, but also other resolvers.

That's my presentation. So for more exploration of how RPKI deployment progressed with the different resolver providers in the last six months, I'd like to defer to the DNSThought website where you can explore yourself. Also on this slide, the research report from Erik and Marius.

I can see I am one and a half minutes overtime.

UNIDENTIFIED MALE: You're fine.

WILLEM TOOROP: I'm fine? Okay.

UNIDENTIFIED MALE: I still don't see any questions for you though.

WILLEM TOOROP: Okay, well….

UNIDENTIFIED MALE: But we're now entering the general question and answer portion of the whole of today's events. So there's essentially 15 minutes, the next 13 minutes for I guess questions for you and for anyone else and anything else, I guess.

WILLEM TOOROP: Okay.

KATHY SCHNITT: You can either type your question in the Q&A pod or you can raise your hand and we will be happy to unmute you so that you can ask your question verbally. It looks like we have a shy crowd for this last session.

RUSS MUNDY: This is terrible. It's a little shocking. How often does the DNSSEC crowd become shy? Very unusual. If we don't have specific questions, just asking from the program committee perspective if there are particular topics of interest that folks would like to hear from in the next workshop either related to things we've had today or any other related topics.

MARK BARROW: Let me ask a question of Willem. That would be with all of this research, where should we be going next? What do we need to do to make [inaudible]?

WILLEM TOOROP: Sorry, Mark, I heard the first part of your question but not the second because the audio was a bit flaky.

MARK BARROW: Sorry. It's going flaky on this side, yes. Where would you like to take us in the future. What will make the Internet [inaudible]?

WILLEM TOOROP: It's getting flaky again. I have lots and lots of things that I would like to see [inaudible] DNSThought. For example, I have looked at this before at the Africa Internet Summit. The latency of RIPE Atlas probes to all the different cloud DNS providers. Is the latency in remote regions or in Africa as good as it is in Amsterdam to Quad1, Quad8, Quad9?

Other things are, of course, DNS over TLS uptake and DNS over HTTPS uptake. We have DNS over QUIC when it emerges, I suppose. DNS cookies, I could monitor the authoritative—no, no, I cannot monitor. I can monitor which resolvers sent them. We could add that to. Yeah, basically anything. Also, if you have an ID, then please send it to me. We can perhaps turn it into a measurement.

RUSS MUNDY: I've got also a question for Willem. That is with respect to the relationship between RPKI and DNSSEC, do you see any way to encourage more uptake of either technology because of the existence of the other?

WILLEM TOOROP: Yes. I don't know but I do know that they strengthen each other. This is because if you have a DNSSEC validating resolver but the domain you're looking into or that you sent a request for a domain which is not DNSSEC signed, then RPKI can still help getting you to reach the correct authoritative name server even though its answer is not DNSSEC signed. So it's improving a little bit. And also the other way

around, if you're resolver is not DNSSEC validating but the domain is signed, then you also do not have DNSSEC and RPKI can help there too. So I think it's a good idea to have as much RPKI everywhere as possible, not only for the end users but also for the authoritative servers and also for the resolver operators.

RUSS MUNDY: Thanks, Willem. One of the things that I have felt would be a very interesting experiment if we could figure out a way to do it would be to determine if there were in fact DNS queries and response that were not signed that were impacted by route hijacks that occurred because there wasn't RPKI in use and to try to collect information that would allow us to correlate because route hijacks happen all the time. I don't know that there's ever been an effort undertaken to try to get some definitization of how many DNS responses were impacted by these hijacks.

WILLEM TOOROP: Yeah, that would be good indeed. But there are a few examples which are publicly known which have a high impact. So I think the message is clear already that it would be a good idea to protect your resolver and the authoritative.

Also, I think especially looking at those PoPs which are not given protection for our beacon, the PoP from CloudFlare, it would be very good if the larger [transit] providers would start RPKI validation too. I know quite a few have already, [inaudible], a few others but that

would make a difference too especially for organizations like CloudFlare that do local peering but still have a default route.

RUSS MUNDY: Mark, have you seen any hands or any further questions?

MARK BARROW: My connection just died and I've just come back again. I didn't see anything.

KATHY SCHNITT: There was no, at this point, no hands raised and no questions in the Q&A pod.

RUSS MUNDY: Well, I think that as Joe pointed out in the chatroom, some of the people may be quite exhausted since the meeting started very early in the day. There are advantages and disadvantages to the virtual meetings, and that's I think one of our disadvantages here. But we do have time for any other thoughts, inputs that you'd like the program committee to hear or take into consideration, suggestions for next time. We're always open to hear more from folks. Anybody care to make any sort of last-minute suggestions? Otherwise, we can close a couple minutes early.

WILLEM TOOROP: I see a question appearing from Hugo Salgado.

RUSS MUNDY: Oh, good.

WILLEM TOOROP: He asks, “Are the RIPE Atlas measurements available from the actual RIPE API?” Yes, they are. All the measurements that are used with DNSThought are listed on the DNSThought website. So you can access them directly and process the data directly. I can maybe copy-paste the URL in the chatroom. Though I have to say for especially the RPKI measurements it’s also better to see the authoritative side, and this will not be seen from the RIPE Atlas measurement results in this case, unfortunately. In RIPE Atlas you only see whether you get an answer or not. But if you want to have access to the measurements from the authoritative sides, then you can contact me and I’m sure we can share something.

HUGO SALGADO: Thank you very much for the URL.

KATHY SCHNITT: Moritz, since you were actually a panelist, you could actually just raise your hand and we can unmute you for you to actually ask a question.

MORITZ MILLER: [inaudible] can you hear me?

KATHY SCHNITT: Yeah, we can hear you. Beautiful.

MORITZ MILLER: Yeah, I was wondering if we should encourage operators of authoritative name server on TLD level but maybe also on second level to make sure that their authoritative name server are located in networks which have assigned ROAs or not. This would be maybe a way forward.

WILLEM TOOROP: Absolutely. The more RPKI there is, the better.

MORITZ MILLER: Is this something that ICANN can do? Should do? I'm not sure. I'm not very familiar with ICANN's structure and what their possibilities are.

RUSS MUNDY: Well, I'm not sure how close that would come to the normal ICANN remit of things. But I can say that as an ICANN entity, the Security and Stability Advisory Committee has a work party underway to address things about routing security. So this is something that as a member of that work party I can make sure it gets introduced into those discussions. Whether or not I'll convince people and it comes out if a document does get published, we'll see. But thanks, Moritz. It's a good suggestion.

MORITZ MILLER: Thanks.

KATHY SCHNITT: Well, it looks like we are right at stop time. I want to thank everyone for joining the DNSSEC and Security Workshop, Part 3, and if you stuck with us through the morning, Parts 1 and 2.

I want to thank all of our wonderful panelists for their great presentations today and managing the Q&A pod like pros along with our moderators who kept us on track with our time. Truly a stellar job by all.

I'd also like to thank the DNSSEC and Security program planning committee for putting together yet another fabulous workshop. These folks work from the end of one meeting to the beginning of the next to ensure success for each workshop.

I want to thank our tech support. Truly amazing folks that make this magic happen for every ICANN meeting, whether we're in person or virtual.

And finally, I want to thank my colleagues Kim Carlson and Andrew McConachie, as if it wasn't for their fantastic teamwork, I would not be able to pull this workshop off.

With that, we will close today's session.

RUSS MUNDY: Before we close, I think it's very important to have especially the program committee recognize the tremendous job that the staff does. Kathy's the one that we hear from, and she's the one that does a lot of the button pushing. But Kim and Andrew are also crucial. Thank you very much to the three of you for the superb job you've done today.

UNIDENTIFIED MALE: Here, here. Always fantastic.

UNIDENTIFIED MALE: Thank you.

KATHY SCHNITT: Thank you, all. Much appreciated. Enjoy the rest of ICANN 69. Thank you.

RUSS MUNDY: Thanks, all. Bye now.

[END OF TRANSCRIPTION]