ICANN70 | Virtual Community Forum – DNSSEC and Security Workshop (1 of 3)
Wednesday, March 24, 2021 – 09:00 to 10:00 EST

KATHY SCHNITT: Thank you. Hello, and welcome to the DNSSEC Security Workshop Part 1 of 3. My name is Kathy and I'm joined by my colleagues, Kimberly Carlson and Andrew McConachie. We are the remote participation managers for this session.

Please note that this session is being recorded and follows the ICANN Expected Standards of Behavior. During this session, questions or comments will only be read aloud if submitted within the Q&A pod. I will read them aloud during the time set by the chair or moderator of this session.

If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, you will be given permission to unmute your microphone. Kindly unmute your microphone at this time to speak.

This session includes automated real-time transcription. Please note this transcript is not official or authoritative. To review the real-time transcription, click on the Closed Caption button in the Zoom toolbar. And with that, I am happy to hand the floor over to Dan York.

DAN YORK: Greetings. Welcome to all of you. I see participants saying they're coming in from all around the world. It's great to see so many of you

here interested in learning more about DNS security, DNS and RPKI and so much more.

My name is Dan York. I'm from the Internet Society. I'm part of the Program Committee that brings you this workshop. We've been doing this for over 10 years now in different forms. And here are the members, some of whom you will be hearing from today. The fact that we have this program is the hard work of all of the folks who are listed on this page to go and bring this together. So when you see one or talk to one, please thank them for their work to make this program possible.

This workshop and the activities here are an organized activity of the ICANN Security and Stability Advisory Committee, otherwise known as SSAC. It is under these auspices which we are able to do this, and so we appreciate that support. There is also additional assistance provided by the Internet Society for some of the infrastructure that makes this possible. So thank you, both organizations, for doing that.

This workshop goes much of the day to day with a great amount of—it's broken into three segments, as you've seen here in the schedule. This part where I will begin has some introductions, some of the status of where we're at with deployment. Then we'll have a panel that Russ Mundy will be moderating, where Moritz Miller will talk about post-quantum cryptography and an interesting part about how that applies to DNSSEC. Scott Hollenbeck from Verisign will be giving us a very detailed look into some of the ways that you should set up a system

and work with things and encryption. It's a great presentation look at that.

Our second piece will be a panel that Steve Crocker will be leading all around automation of provisioning and the various different technologies, where we're at with the different pieces, what's going on with that. It's a critical part as we look at how do we automate more of DNSSEC and make it just more resilient in different ways.

The third part of our session at the end will be moderated by Fred Baker. As a number of different sessions talking about measurements, Ed Lewis from ICANN will be here talking about RPKI, Route Origin Authorization deployment, visualizations of the DNSSEC. Victor has some information on NSEC3. Wes Hardaker has some other information around the deployment of a piece. And so it'll be a good session in that part at the end as well.

So that's what we're going to look at over the scope of this time. I hope you can join us for as much of the presentation as you can.

As Kathy mentioned at the beginning, it will be recorded so if you don't get a chance to listen to directly, you can go back and listen to it. Throughout the session, as Kathy also mentioned, please do feel free to post questions through the Question and Answer pod. We will get to them as we can.

For my first part here, I want to talk a little bit about where we're going with what trends we're seeing overall. The first is we look at the two sides of DNSSEC are around the signing of domains—and we'll talk

more about that—and also the validation, the checking of those signatures. And both of those are important for the ongoing piece.

The tool we've been using to watch validation for quite some time has been the stats that the Asia-Pacific NIC (APNIC) operates. And on their charts they continue to show that on their global statistics, about 25% of all DNS queries are being validated or coming from networks where validation takes place. That's really what it is. 25% coming from those kind of places. They go into more details and we can see a number of different areas and regions where we go much higher. If you go to their stats, you can actually dive deeper in and see what specific country or region may be going down into the specific networks, and seeing which networks are doing the most for DNSSEC validation in different kinds of ways.

The DS record is another mechanism that we're looking at, which is how many of these—these are the signatures that are transmitted from the registrars up into the registries, and so we're up at this top level there. And we're seeing this count, which if we look at that, it's over 14 million coming up into that space around globally across the TLDs that are being tracked by the DNSSEC-Tools project. And you can go to that site, stats.dnssec-tools, and be able to see specifically what's noticing there.

We also look a lot at the use of DNSSEC to sign MX records and provide DANE records to allow for the secure sending of e-mail from server to server using DANE as a mechanism to go and retrieve the TLS certificates and be able to go and have secure encrypted e-mail. We've

seen an ongoing rise as this chart shows over the last couple of years. Viktor Dukhovni, who you'll hear from in section three, is one of the ones behind this tracking in the pieces, as well as Wes is also involved with this too. You can ask questions to them if you want to know specifically about some of these measurement components. But you can see we're up around two and a half million domains here that are currently being tracked, and this is a wonderful continued growth that we're seeing here.

In this talk at the end, especially, we'll talk a little bit more about RPKI, the Resource Public Key Infrastructure, which is a key component in how we secure the routes on the routing system that underpins the Internet as far as getting traffic from one place to another. If you're not familiar with what goes on within RPKI but basically routers, things, send out/distribute what are called Route Origin Authorization. It's basically saying, "I have the authority to go and originate this route." And we're seeing a nice growth in this, as you can see along here, and the percentage of the unique prefixes that are being shown. And so it's good to see this growth. It's something that needs to happen to secure the routing layer of the Internet, and you can learn more at the site that's mentioned here, which is operated by NIST there in the United States.

We also have some good stats coming out about what each of the Regional Internet Registries or RIRs, what the growth is of the number of ROAs in each of these of the RIRs. And you can see this nice curve as we've tracked up with a lot of continued growth in RIPE, some more

recent growth in ARIN, and continuing on from there that we see this growth happening. It's great to see this kind of thing going on.

For a good number of years now, we have also been maintaining a set of maps that show the deployment status of country code TLDs (ccTLDs). Now, obviously, we're also interested in the growth of DNSSEC signed domains in generic TLDs and also in all of the new generic TLDs as well, but those can't be mapped. You can't easily. There's no sense of that. So geographically, we look at what the ccTLDs are doing.

Over the time, we're now up to where we're getting a significant number. We have over 137 of the ccTLDs have a DS in the root. They've signed the top-level domain and they're accepting domains in different ways. So this is wonderful to see, tremendous growth of what's going on. If you're interested, you can sign up to get these maps delivered to you every Monday morning and you can see the latest statistics that are there.

One of the things we found is that over time these particular maps have gotten more and more green, green being the top state in the current mechanism. So one of the questions we've asked is, how can we evolve these maps a bit more to show the next desired state?

So right now we are currently tracking five states: experimental, announced, partial, DS in root, and operational. As we've seen this, we now have 137 of the countries are in these last two states, DS in root and operational. And so the maps, in particular in Europe, in North America, etc., are filling in and they're becoming all green, which is

great, success, but we want to use the maps to help us see where do we want to go next.

So in that second session of our work today, you'll be hearing all about DS automation and how we go and automate the provisioning of this. So with that, we're going to be evolving the maps to add a sixth state which will be DS automation. So do the top-level domains in the ccTLDs use one of the various mechanisms that we have to go and do that? A lot of what we're seeing is pulling for CDNS records or CDNSKEY records, things like this.

And so, where we're going is that the maps will start to look like this. This shows a couple of the domains, .ch, .cz, and .sk that are currently providing DS automation. So this is what the maps will be looking like. This is starting now, and so you will see these maps start to show this in the time going forward.

So if you are with a ccTLD and you do implement DS automation, I would ask you to please send me a message, york@isoc.org, and I can get you updated in the database so you can show up in there. And if you subscribe to the maps, you will start to see them now coming out with this in progress.

So with that, I will just say that we do have various resources out there that we would encourage people to look at to learn more about how to deploy on DNSSEC. DNSSEC-Tools, I've mentioned the Stats site. There's also a great amount of information around there. The Internet Society ran a project last year called Open Standards Everywhere that has a number of different documents up in GitHub that you can go and

look at to learn about how to deploy DNSSEC. There is an older site, dnssec-deployment.org, that has a great amount of historical information. And we already talked about the stats that APNIC provides. There's also a good number of resources for RPKI, and we will be talking a bit about some of that a little bit later in the presentation as well.

So with that, I am going to say thank you very much. Again, welcome to the session. I hope you enjoy the data we have ahead of you. Kathy, unless there are any questions, I will turn it over to Russ. But I'll first ask, are there any questions right now?

KATHY SCHNITT: I don't see any questions at the moment.

DAN YORK: Well, excellent. Then I will say thank you very much, and have a great day. I'll turn it over to you, Russ.

KATHY SCHNITT: Thank you, Dan. Russ, all on you.

RUSS MUNDY: Okay. Thanks very much, Dan and Kathy. I appreciate that. This is our first panel of today. Our first presenter is going to talk about an area that we have not dealt with directly in workshops earlier in time though there's certainly has been a number of discussions around the

community. And it's the impact of some of the results or implications of post-quantum computing on DNSSEC and other related activity. So, Moritz, why don't you please go ahead? And now I'll turn the floor over to you.

MORITZ MILLER: Thanks, Russ. So this is actually a joint presentation with my colleague, Jins. I hope he's on the call as well. Jins, would you mind sharing this screen?

JINS DE JONG: I have tried but I do not have the permission yet.

KATHY SCHNITT: Tech folks, can you please give Jins co-host rights?

MORITZ MILLER: So this is research that we've carried out together with folks from TNO, which is a Dutch research organization—again, I work for the Dutch ccTLD—folks from NLnet Labs, and also University of Twente. We have given this presentation in different forms at different venues already. But today we want to give you a quick update on the NIST competition, which has to go to standardize post-quantum crypto algorithms and discuss briefly the impact of these algorithms on DNSSEC. And we also want to take the feedback that we got from previous presentations and take them here as well to discuss with the

community as well and refer the ones to give you an outlook on the research that's still yet to come. Next slide, please.

So, just a brief introduction. You might have heard that there are these quantum computers that are being developed at the moment, and these have the potential to break current public-key cryptography. And this would also then affect DNSSEC because all the algorithms currently used in DNSSEC could be potentially broken by these quantum computers.

Luckily, new quantum-safe algorithms are being assessed at the moment by the standardization organization NIST, and we wanted to understand whether these quantum-safe algorithms are suitable for DNSSEC or not. But first, Jins will give you an introduction to post-quantum crypto and also on the most recent developments.

JINS DE JONG: Thank you, Moritz. As Moritz already mentioned, some of you may have seen this presentation in a slightly different from before. Therefore, we've added some new developments and ideas to make it interesting for those as well. I will now focus a bit more on the cryptographic aspects. And afterwards, Moritz will present to you the implementations for DNSSEC.

The relevant thing here is that quantum computers may appear and there is an algorithm known to run on a large quantum computer known as Shor's algorithm that breaks all public-key cryptography. Public-key cryptography is used in DNSSEC for signatures. So that

means that when this large quantum computer appears, DNSSEC's current signatures are no longer safe/secure.

When may this happen? This is still quite a while ahead of us, but last year some colleagues of mine made an estimation and the earliest possible would be in the 2030s. But of course, if science has a bad day, it could take a few decades longer as well.

Why didn't worry already about this? The thing to understand here is that we have plenty of time but we also need plenty of time. This is what practically explained by Mosca's inequality, which on one hand says, "As long as you need less time than you have to prepare for the quantum computer, you're okay." And how can we define the time we need? That's, on the one hand, the time we want our secrets to remain secret. In the case of DNSSEC, this would be the secret keys to generate signatures. But these can be changed fairly quickly.

What is hard, however, is the other part, the time to switch to another mechanism or to introduce a new signature. This could easily take a decade. Well, the earliest estimation is that within 15 years such a quantum computer could exist, it's time to consider slowly moving towards quantum-safe signatures.

Precisely because of this reasoning, NIST has started a post-quantum cryptographic standardization competition in 2016. They haven't done this because there is already a perfect candidate and the outcome was clear in advance. It was much more because alternatives to the current public-key cryptography are needed. This can also be seen in the outcome of especially the first round, where many ambitious

researchers submitted their proposals and several of them were broken within months, many more than would have been so in other competitions.

We're currently in the final round, which is expected to last until the end of this year. And from the signatures, there are three signatures as finalist and three other signature schemes are mentioned as alternative candidates that may be standardized now or in the future and mainly serve as an alternative to be aware of the final.

These are remaining algorithms and what is interesting here to note is that they're very different from what is currently used. Currently in DNSSEC, one of the algorithms is the elliptic curve down at the bottom. We have shown some performance indicators in this table. We compare them to the various finalists, alternate candidates that are still in this competition, and they're wildly different.

Private keys are very different but that's not so exciting from a DNSSEC point of view. What is relevant are the large or very small public keys, and especially the signatures. Is there among these candidates a signature scheme that is suitable for DNSSEC? That is what Moritz will discuss in the second part of this presentation. For now, it's mainly to see how different they are.

Also since the start of the final round last July, there have been some developments. One is that there have been new attacks and improved attacks and a better understanding of the security of the multivariate algorithms which, on one hand, increase our knowledge of them, but on the other hand, also show that we do not have great trust in them

yet. That is still something to consider that one day, someone may come up with an attack that breaks them or of course any of the other schemes. Another issue is that one of the finalists, Rainbow, does not yet offer a royalty-free implementation.

These issues together have brought NIST to express some concern that there may not be enough diversity among the signature scheme's candidates, the candidates for a signature scheme, so that certain difficult use cases may not find a suitable candidate. This has led to some ideas, could we do something else? And I didn't have complete time to explain this idea in detail. Therefore, we have put the link to the original idea in it from Verisign. But the idea is to use a construction of a hash-based signature for DNSSEC as well. Where at the bottom, there would be records. And consecutive hashes would finally yield a top hash, which could then be used as a public key. And then only the path towards the top hash would demonstrate authentication. However, to not nibble on Moritz's time, I'll give the presentation to Moritz.

MORITZ MILLER: Thank you, Jins. Let me also briefly share the screen. Jins pointed out already we have seen that this new discussed algorithms do have some quite different attributes. So we were wondering would that have caused any issues when we would have tried to apply them to DNSSEC. Just understand that we tried to identify which limitations/restrictions does the DNSSEC, DNS and the underlying transport protocols have. For today, I only want to focus on the key

and signature sizes. Because here, we've seen from previous research and our own measurements that DNS messages larger than 1230 bytes quite often cause fragmentation. People argue also that this number could be a bit bigger, but around 1230 bytes, we do sometimes see troubles with transporting DNS messages.

Also, DNS in general is quite attractive for DDoS attacks and it has been misused in the past very often. So that we want to make sure that we don't have too big of records, DNS messages and sequences of keys, such that we make DNSSEC not even more attractive for this kind of attacks.

If we then look at the algorithms that are still in the third round of this competition, and if we only look at algorithms that have signatures below this threshold of 1230 bytes, then we already see that at least Falcon on first sight seems to be a quite adequate algorithm for DNSSEC. However, as soon as we would like to transfer multiple signatures or multiple keys into one single message, then we already might run into troubles, because then we already have DNS messages larger than 1230 bytes.

The signing performance and verification performance seemed to be all right for the algorithms. You can find more details in the paper that I've linked at the end of our presentation.

If we then look at two other algorithms in one of the NIST finalists, we see that we have Rainbow-1a and RedGeMSS128, which both have great signatures. They're very small signatures which are on par with signatures that we've seen with ED24519, for example. Their signing

performance and verification performance is good, at least from the first measurement that we're doing. Unfortunately, the public key is very, very big. It is even that big that it can't fit into regular DNS messages anymore because there we have a limit of 64 kilobytes. For this reason, we see two main challenges. The first are keys above 1230 bytes, and the even bigger challenge is we might get keys which are larger than 64 kilobytes.

The first issue could be addressed relatively easy. DNS already has a solution for that, which is TCP. TCP is just regular DNS. It has already implemented all of the regular DNS software. However, it might not be everywhere supported. We do see that sometimes middle boxes do not like DNS traffic over TCP and block it. Or we see that people think that DNS is UDP protocol so they don't support TCP.

Also, if we would start sending all the same messages across TCP, we might see an increase in several requirements. However, if you look at the last two points that there has been researched that this might not be completely true. So in the first case, we've seen that DNS, some servers do not support TCP. That is definitely the case. But the number is quite low, below 1% according to this blog post that I've linked to below.

Also an earlier study from 2015 has shown that sending DNS messages via TCP and even including TLS might not come at a too much decrease in performance. So sending everything via TCP might be an option in the future.

This leaves us with a second challenge, the keys larger than 64 kilobytes. Here we propose two possible solutions. The first solution would be splitting the key into multiple resource records. This would be, in our opinion, a modest DNS extension but has the disadvantage that we would require additional round trip times to transmit all the additional research records. And it might also lead to a higher chance of packet loss. Because every time we have to send a packet, we might also increase the chance of packet loss as well.

Another solution would be to distribute keys out of band. This mean that we would introduce a new DNSKEY record, which then includes a pointer towards a different location, for example, on that server, where then the recursive resolver could go there and fetch the key using a different protocol like HTTPS. This will be less prone to packet loss but it requires support for different protocol, and also then the DNSSEC operators would have to provide additional service to provide the key.

Both of these solutions have somewhat the advantage that in general, keys are not exchanged very often. In most cases, a TTL of keys are one hour or even longer. So we don't have to rely on these mechanisms too often. But of course, both solutions add to the DNS Camel. So this means that we would have to extend the DNS even further, which not everyone likes. And as we got the reaction at an earlier presentation at DNS Org that people suggested that we just start over again and start with a new version of DNS so that we don't have to work around to these issues anymore.

To conclude this short presentation, we believe that we can apply quantitative algorithms to DNSSEC, but we also think that will require some modification. Also, we've seen that things are still in development, so things might change, and so we might have to think outside the box and consider, for example, these hash-based style algorithms in the future.

The next step of our research is to also simulate the impact of post-quantum crypto on real DNS traffic. For example, imagine simulating what happens if you could replace all the signatures that we see as a big resolver with a post-quantum crypto algorithms and what would be the impact of that and also implement some of the proposed solutions as well.

Also as Jins mentioned in the beginning, rolling to a new algorithm does take time. I've presented earlier research at the last DNSSEC workshop about this as well. And this is why we think that we should think about transitioning to post-quantum crypto algorithms as early as possible. On this slide, you can also find the link to the paper. With that, I would like to thank you for your attention.

RUSS MUNDY: Thank you very much for that presentation. One of the things that I wanted to mention is that SSAC did issue a very short publication, but it was directed to the NIST review panel that is ongoing. That was in December of 2019. So one of our sponsoring organizations has been watching and taking interest in this, but this is a really good study and goes into a whole lot more depth than what the SSAC paper did.

We have a very good 15-minute session of Q&A at the end, and I'm not seeing anything. I do one thing in the pod. Okay. This one is from a Daniel Migault. Daniel, this is very much specific to your presentation. So why don't we go ahead and take this and then—can you see the Q&A pod, Moritz?

MORITZ SCHNITT: Yes, I can.

RUSS MUNDY: Okay. Would you go ahead and take or attempt to answer Daniel's question, please?

MORITZ SCHNITT: Yeah. Let me read out the question first. Daniel says, "One aspect is the size of the key-in signature of one specific post-quantum algorithm. In the case of multiple post-quantum algorithms being provided in a given zone, I'm wondering how likely it is that the client specifies one of these algorithms be returned only the information associated with the key? Typically, it seems to me that the [inaudible] DNSKEY version all possible keys. Is this an issue? And if so, how do we intend to deal with this?"

I think the current form of DNS protocol, just every single key will be returned. The resolver doesn't have any means to signal which key it would like to receive. As I remember, I think there was an academic

proposal a few years ago to start with some kind of not key exchange but key agreement protocol, which algorithms should be used. So maybe we might want to revisit that at some point in time so that we can reduce the number of keys that we have to transmit. If I can find the paper, I will link it to the chat.

RUSS MUNDY: I believe I remember that paper also. But you've probably looked at it much more recently. That was, I think, a good answer.

Let's see. Okay. I think at this point, I don't see any more specific questions for this presentation. So let's go ahead on to Scott. And we'll have, like I say, a few minutes up to—I think we have 15 scheduled but we have until the top of the hour for this session. Scott, over to you.

SCOTT HOLLENBACK: Thank you, Russ. Thank you very much for having me today. My name is Scott Hollenback, Verisign fellow. The title of this presentation is "A Balanced DNS Information Protection Strategy: Minimize at Root and TLD, Encrypt When Needed Elsewhere." It's based on a paper of the same name that was authored by Verisign Senior Vice President and CTO Burt Kaliski, and previously published on the Verisign and CircleID websites. Next slide, please. Thank you.

The DNS is in a new era of change with several proposals that include an increased focus on confidentiality protections being discussed in places like the IETF. Examples include data minimization techniques, such as QNAME minimization, and cryptographic techniques such as

DNS-over-HTTPS, sometimes known as DoH, and DNS-over-TLS, sometimes described as DoT. These different approaches are designed to provide confidentiality protection for the different exchanges of information that take place during the DNS resolution process. For the next few minutes, I'm going to describe Verisign's current recommendation as a root and TLD authoritative name server operator to minimize at root and TLD, encrypt when needed elsewhere. Next slide, please.

So how did Verisign develop this recommendation? We believe that there needs to be a balance between the elements of the confidentiality, integrity, and availability or CIA triad. DNSSEC design, for example, factors in operational risk. Otherwise, all name servers would be expected to have ZSKs that would be signing zones in real time. Protocol changes such as DNS encryption can create new operational challenges, expand the attack surface, and impair the ability of network operators to manage their networks. Additional layers of software add complexity as well. Complexity adds fragility and fragility can lead to outages.

Consider that NIST National Vulnerability Database has recorded 950 vulnerabilities associated with TLS since 1999, and 347 vulnerabilities over the last three years. New and unmitigated vulnerabilities can have an impact on server availability if exploited. So the operational risk of adding any new feature to the DNS camel has to be balanced against the disclosure risk that's being addressed. The risk benefit trade-off is different for different exchanges in the DNS ecosystem. So

we shouldn't consider this a one-size-fits-all recommendation. Let's take a look at the different exchanges. Next slide, please.

First, the client to recursive resolver exchange. Here, the resolver sees client specific information and full query names. If a single resolver is used, that resolver sees all of the domain names for which resolution has been requested by the client. A passive observer can see the same things. Next slide, please.

For these reasons, we believe it's appropriate for clients and resolvers to implement DNS encryption to provide confidentiality assuming that no other cryptographic protection such as on a network connection is available. Remindful that there may be enterprise network management concerns to address when adding encryption to a traditionally unencrypted network protocol. Next slide, please.

Now let's look at the exchange of information between the resolver and the root and TLD name servers. Recursive resolvers represent the aggregate interests of their clients such that client-specific information is not sent to the root and TLD name servers. However, with traditional DNS resolution, the resolver sends the full query name such as www.example.tld to the root and TLD name servers, which is more than what they need to know to perform their resolution task. If the resolver implements QNAME minimization, the name servers will see only the aggregate interests in the top and second level domain names that appear in the queries. So what form of confidentiality protection is appropriate for this exchange? Next slide, please.

We believe that minimization techniques are the best option for this exchange. Disclosure risk to both insiders, which include the root and TLD name servers, and outsiders, such as passive monitors, has been significantly reduced. With QNAME minimization, the root name server sees only the TLD for which resolution is requested. The TLD name server sees only the second level domain such as example.tld. Meanwhile, name server availability affects navigation to the hierarchy below any particular level of the DNS. As such, a root or TLD name server outage can have an impact on millions of zones. So it's not clear at this point that it's a good trade-off to focus efforts on adding encryption to the exchanges, given that QNAME minimization and similar techniques already reduce the risk of disclosure of sensitive information. It's better to focus elsewhere, at least for now. Next slide, please.

So how widely deployed is QNAME minimization? It's been implemented in many recursive resolver implementations and its use is growing impressively. In February 2021, 55% of the queries that Verisign observed at the .com and .net name servers were minimized. And that's up from 32% of the queries that we observed in January 2018. So the trend is positive and increasing. Next slide, please.

The next exchange of information takes place between the resolver and the second level and deeper name servers. This exchange continues to convey the aggregate interests of the resolver's clients, but now they may include full query names and client-related information such as the client's subnet information to improve

resolution performance. The disclosure and operational risks are more balanced for this exchange. Next slide, please.

As such, we believe that encryption gives an appropriate risk benefit trade-off between resolvers and second level name servers when sensitive, full domain names or client-specific information needs to be protected. Next slide, please.

So let's look at these different techniques again. DNS confidentiality protection is available using two different types of techniques: encryption and minimization. Encryption uses cryptography to conceal information, reducing the risk of disclosure to outside parties. Both participants in an exchange must support the use of encryption and there is the potential for operational impact on both participants. DNS-over-TLS (DoT) and DNS-over-HTTPS (DoH) are both examples of encryption techniques. We encourage measurement studies that explore how well these techniques stand up to attack scenarios. Minimization techniques decrease the sensitivity of information by reducing what's exchanged to only what's necessary to perform a requested function. This reduces the risk of disclosure to both outside and inside parties with no operational impact on the receiver. Query name minimization, NXDOMAIN cut processing, and aggressive DNSSEC caching are all examples of minimization techniques.

And this leads us to—next slide please—our conclusion that DNS encryption and various minimization techniques all have a place in protecting different DNS exchanges. We believe that it's best to balance the risks of disclosure and risks to availability depending on

the nature of the information being exchanged. As such, our recommendation for DNS confidentiality for now is to minimize that root and TLD and encrypt when needed elsewhere. That's my last data slide.

RUSS MUNDY: Okay. Scott, thank you very much for that presentation. And I am looking—we have a question in the Q&A pod. Oops, let me get to the right tab here. We have a couple questions in the Q&A pod. It looks like, from what I'm seeing, these are directed to you, Scott. Can you see the Q&A pod and take a take a walk through them?

SCOTT HOLLENBACK: Yes, I sure can. So let's start with Peter, "Do I imagine Verisign's position shifting when sufficient operational experience has been gained in the SLDs and below? In other words, do you see .com, .net, and the root deploying encryption at some point in the future?"

Well, the future can be quite far off. Right now, we are not currently planning on deploying any kind of encryption support on the TLD or root name servers that we operate. We're still very concerned about understanding the operational risks, the impact of outages, etc. We are encouraging experimentation, though, at lower levels of the hierarchy. And we believe that unlike DNSSEC, where success largely depended on us doing things from the root and down, success to deploy these technologies will probably depend on the deployment success at the leaf nodes of the tree and up. If we can demonstrate

that these techniques are safe and don't increase the risk of outage or other impacts, it's something to consider in the future.

It's also possible that different techniques—right now we're talking about DoH and DoT, but the IETF is also looking at DNS-over-QUIC. QUIC tends to be somewhat less resource intense than TCP and TLS, and so it might be a viable alternative in the future. However, it's also a very new transport protocol so we don't have a lot of implementation experience, we don't necessarily have native support in operating systems. So it'll be interesting to see where that goes in the future. That's Peter Van Dyke's question.

Now, Peter Thomassen. Let me see if I can answer this one live. "Drawing the border at the SLV server is not always correct. For example, in the case of .co.uk, the border should probably really be at the registerable domain versus public suffix. Have you considered using the public suffix list to allow resolvers to decide whether to minimize?"

Well, of course, that's not really our decision, Peter, but you're absolutely correct. The presentation is looking at this from Verisign's perspective where those lines are more clear. But when you start talking about zone cuts and how different name servers actually support zones across cuts, yes, that's definitely something to consider.

All right, more questions moving down. Question from Nicolas. It looks like it's more focused about quantum and crypto. Potential issue of increasing signature length... I'll leave that one to Nicolas.

Viktor Dukhovni: “Isn’t the second level domain often the entire client queries suggesting that TLD traffic is often sensitive?”

I think that’s correct and that’s why we’re recommending minimization, Viktor. I’m sorry. In the case of this, no, you’re right. The second level domain. If the entire query is the full QNAME, that’s where we’re recommending encryption.

All right. Then the last one I see directed to me is Vittorio Bertola. “Do you think that each TLD root operator should decide on their own, whether to support encryption or would you prefer to have some agreed uniform policy? If so, is this a job for ICANN or for the IETF or for where?”

Well, I do think that there’s a policy element to this. And this is a question that groups like RSSAC and SSAC—they’re going to have to look at this. The IETF being protocol specification people, for the most part, I think it’s wise to point out security considerations, operational impacts, but I don’t see the IETF being the entity to produce any kind of policy statement on where or how this might be appropriate. And I do believe that this is a place where TLD and root operators may decide to experiment on their own. If you’re operating a very busy TLD like .com or .net, where you’re processing millions of queries per second, your concerns are going to be very different than if you’re processing a TLD that’s processing a much smaller amount of traffic and the query loads are smaller. Those are all the ones I see addressed to me. Thank you.

RUSS MUNDY: Thank you, Scott. I appreciate that. Now, if we could go back Moritz. Could you pick up the ones that look like they're addressed to the quantum space? I promise not to move it again like I did on you before. Sorry about that. There we go.

MORITZ MILLER: Sure. Let me give it a try. Let's go first from Nicolas, "Scott's talking about quantum encryption and to encrypt when needed. It seems this potential issue of increasing signature length will may happen many times during a regular webpage retrieval, for example, UX, DNSSEC, SSL, etc. Is there any initiative to try to kind of share some of these protocols among some of or all protocols so as to minimize the data exchange increase so it shares some of those credentials?"

As far as I'm aware, no. I'm also not sure how this should look like. We have DANE so we could probably put some keys in there as well in the DNSSEC. But I'm not aware of any of those initiatives. I'm not sure how this should work.

RUSS MUNDY: If I could add just a little bit here, Moritz. That is from a historic perspective, it has been incredibly difficult to try to have a protocol mechanism used for security purposes between various other protocols. There are a couple of examples where this has been done in some of the IETF specifications, but not very many. So it's a very interesting idea and I'd encourage folks that think we ought to go this way to start giving it more thought and looking for how such an

approach might be accomplished. Thanks. Go ahead and go to the next one, please.

MORITZ MILLER: Thanks, Russ. So next question from Sivasubramanian. Sorry for butchering your name, "Out of the threats to the current DNS signature algorithms by quantum computers, is there an interim problem due to the present changes and present advances in desktop computing architecture, where in parallel computing technologies are becoming more and more real and even otherwise far more powerful processes are making their way to the desktop workstation computers?"

If I understand the question right, she's wondering about whether advances in current computers might threaten the security of DNSSEC at some point. I think all the algorithms are already always under threat by developments and by security researchers trying to break them. So I think I don't see any of the developments that the one who's asking the question is pointing out as a current threat, but I think we always have to be ready in DNSSEC and other security protocols to exchange the algorithms at some point, regardless of whether there's quantum computing or not.

RUSS MUNDY: Okay. Thanks. I think the next one from Hugo is a more general question, one that has been raised before in this workshop. But I'd be interested in hearing Moritz's take on where do you see DNS signature

and validation occurring. I have my opinion, but I think a lot of people have already heard that. So I'd love to hear yours.

MORITZ MILLER: I think, personally, preferably the clients because then we would have the actual end-to-end security there. I'm not sure if this will happen in the future. At the moment, it's at the resource. That's key. And they're signing also everywhere. I think we are still far away from that.

RUSS MUNDY: Okay. Thanks. I guess the next one is a follow on.

MORITZ MILLER: I think this is a clarification. The last one from Nicolas, "What is the current state of the cryptographic algorithms in other protocols such as RPKI and BGPSec, in reference to post-quantum encryption, and also taking into account that they maintain a similar model based on trust chain?"

Honestly, I'm not aware of any research regarding post-quantum crypto on RPKI, for example. There is, of course, research on the TLS field, quite a lot of research going on there. But I'm not aware of research on RPKI, for example. I'm not sure if Jins knows about any.

JINS DE JONG: No, I do not. I'm sorry.

MORITZ MILLER: Okay. And then Viktor Dukhovni is asking, "Since DNSSEC's only concerned with confidentiality, we can probably wait longer than other protocols that do post-quantum crypto algorithms."

Jins, do you want to take this question?

JINS DE JONG: I think that's a fair point. DNSSEC is not the most threatened system to be attacked by quantum computer. Nonetheless, if possible, to keep it secure, I think once you try to do so and start the transition early. Laziness should not be the reason not to guarantee secure DNSSEC.

RUSS MUNDY: We do have a comment from Suzanne Woolf in the chat room. "For what it's worth, it's entirely possible for the IETF to publish best practices, RFCs. And DNS has published advice to operators on various things." But as Scott says, "We need operational experience and encrypted DNS techniques first, and that will take time." I think that's primarily related to Scott's presentation. But I think it also relates very much to the post-quantum crypto and particular in the aspects of algorithm rollover. This is an area that indeed does need further study, in my view, and I think a lot of people agree that we need to do more work in the space of what it takes to add or decrement an algorithm from DNSSEC.

We still have a few minutes for questions if anyone has any more. Okay. We just had another one in our Q&A pod. Okay. "Does the

Internet importance require a uniformity of technical developments regarding ccTLD?"

I'm not sure that I understand the question. Do any of the other panelists see and want to take a—perhaps Clement, if you could add just a little more clarification or we could activate your mic and you could ask it verbally.

JACQUES LATOUR: Hello, Russ. Jacques here.

RUSS MUNDY: Hi, Jacques. Go ahead.

JACQUES LATOUR: I think the question is, do all the ccTLD need to implement the same services across the board? Do we need to all respond to quantum or all do QNAME minimization equally?

RUSS MUNDY: Okay, good. It sounds like you're in a good position to answer that, Jacques. So if you would add some response more, please do.

JACQUES LATOUR: I guess we need some level of uniformity across all the ccTLDs, but there's no requirement for all of us to be at the same level. So I think that's the answer.

I did have a question for the quantum. Assuming we'll never know when people will have the ability to control quantum to do some damage or we'll never know when people are actually able to break the protocol. So I think it's best practice for us to implement as soon as possible quantum protocols before it's too late. So we have an opportunity to be proactive. I guess the question is, when do we need to be proactive on this?

MORITZ MILLER: This is a very hard question. I think by bringing this topic to venues like ICANN and IETF, we take the first important steps. And I know that also other people are very interested in this topic, then, hopefully, be ready when we need to. I think it's really important to start thinking about transitioning to these algorithms as soon as possible. I think I can't give you any more details on when we should do that. I think that's a one million dollar question.

JACQUES LATOUR: Thank you.

MORITZ MILLER: But if you're interested in discussing this topic, then feel free to reach out to us.

RUSS MUNDY: So we do have a follow-on question or a comment from Clement. I'm reading the entire question, it's really reflect some of the work he did

earlier in evaluating or looking from a registrar's view of things and he's asking if ICANN should update technical minimums for some ccTLDs. The relationship between ICANN and ccTLDs has always been one that's had an interesting perspective from both sides. I will see if Jacques would speak to this as a person who is closely associated with a ccTLD and the group that does work in the ICANN space on this. You're still muted, Jacques.

JACQUES LATOUR: I don't have more to say than I said before. It's up to ccTLD to implement the best practice. That's why we're all talking here today.

RUSS MUNDY: It really is a cooperative kind of arrangement. There's no real strong directiveness that ICANN really exerts over the ccTLDs because everybody is operating them for the best good of the Internet.

Let's see. Moritz or Jins, could you read Nicolas's question, please?

JINS DE JONG: Should I take this one? "Have you considered in some research to include quantum key distribution as a DNS security enhancement? Considering current prepared infrastructure in Europe for QKD example for exchanging purposes, my understanding is that this would ensure a secure exchange anyway with current weak algorithm keys."

Well, the answer about the first question is we have not considered it. We have started with the standardization efforts by NIST and assume that the most likely candidate would be a standardized algorithm, which was the reason to look at this competition. Only when in this process, it became clear that many of the candidates do have some complications/issues and are not trivial to implement in DNSSEC, we have asked ourselves and Google the question, what alternatives would there be?

As far as I'm aware, most of the QKD proposals have problems guaranteeing a sufficient bit rate to set up a key yet. Perhaps this may be an option in the not-so-far future. I wouldn't dare to say so.

RUSS MUNDY: I see we have reached the top of the hour. I think Viktor's comment is really more of a comment than a question. Pablo's question—Moritz, can you do a quick response to Pablo's question so we can let folks get to the break? And then Daniel's as the closing question.

MORITZ MILLER: Sure. Pablo, I don't know this paper very well myself, but I would like to point you to a proposal in the IETF and that was also pointed out by Shumon in the chat. So maybe you find more details on why this proposal didn't really take off. Or maybe Shumon can answer the question. But I will post again the link to the draft that Shumon post in the Q&A as well.

RUSS MUNDY: Okay. Daniel Migault's question, "I'm wondering how post-quantum can be experimented. As when published, everyone will retrieve all post-quantum and non post-quantum responses." Moritz or Jins, could you take a quick shot at answering that?

MORITZ MILLER: I'm not sure if I completely understand the question. But I assume that is how can we experiment with post-quantum algorithms if, for example, the root has not deployed it yet? I think our first step would be to take recorded traffic and simulate what would happen if we would replace all the signatures and keys that are being transmitted with the post-quantum crypto algorithms and then see the effects, if this could be one of the first stages.

RUSS MUNDY: Okay. Thank you very much. Thanks very much particularly to our panelists today and for all of the questions that we had from the workshop participants. This has been a most addressing session from my perspective. I hope others found it useful. I'm very glad we've got some good feedback in the chat room. Sorry, we got a little bit long but it looked like we were having a good set of questions. So I will apologize for consuming a little bit of the break time. But with that, we'll close the panel, and back to Kathy for the end remarks for part one.

KATHY SCHNITT: Thank you very much to all our panelists and moderators. It was a great session. We will meet back here at 15:30 UTC for part two of our workshop. We are in the same webinar room so there is no need to disconnect. You are free to just stay and hang out with us and listen to us chit-chat on our side. Thank you. Please stop the recording.

[END OF TRANSCRIPTION]