BARCELONA – DNSSEC Workshop (3 of 3)
Wednesday, October 24, 2018 – 13:30 to 15:00 CEST
ICANN63 | Barcelona, Spain

UNIDENTIFIED FEMALE: October 24th, 2018. DNSSEC Workshop (3 of 3), 1:30 – 3:00, Room 113.

RUSS MUNDY: Okay, DNSSEC Workshop folks, we are ready to gather up and restart here momentarily.

Well, I think the first thing we need to do is to say thank you very, very much to our lunch sponsors. I hope everybody had a nice lunch. I know I enjoyed it. It's Afilias, .ca, and SIDN. Thank you very much.

JACQUES LATOUR: Alright. Welcome. So, we're going to have a panel for the next hour-and-a-half, almost, a presentation and panel discussion on DANE. So, first up is Wes Hardaker.

You're going to do Viktor's presentation, I guess? Or, yours? A mix?

WES HARDAKER: Yeah. It's a combination of both. So, most of this material did come from Viktor Dukhovni, who's the person on the left-hand side of this slide. I'm Wes Hardaker from USC-ISI. This will be – in fact, even the title is slightly wrong because we added a bunch of stuff – DANE and SMTP

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

in particular and the usage thereof, as well as some general DNSSEC statistics.

So, we'll go over a bunch of things. I'll start with a little bit of background about SMTP works. We'll talk about e-mail security without DANE and what that looks like and what it looks like with DANE and why it's better, and then we'll look at some DNSSEC and DANE deployment statistics that Viktor has gathered.

So, first off … oh. I'm looking at my slides. That doesn't help. There's the overview slide. Can we zoom out a little bit?

Thank you. I need to be able to see that diagram. Even more.

Okay. You don't need to read the words. I'll talk to the words. Yeah. Thank you.

So, when you send e-mail and when you read e-mail, you probably know that you're using secure stuff. Your laptop sends, on the left – some user will send to their ISP, and they send it over. You need your user name and password, and it goes over TLS-protected SMTP to your ISP.

On the right-hand side, you have IMAP. IMAP also requires TLS. You connect to your ISP to pull down mail.

The problem is that, between these two points, as Viktor puts it, a miracle occurs because, between the servers, it actually is unencrypted and unauthenticated, which is just sort of fascinating in this day and age.

So, if you had submitted this as your architecture to some class project in a computer science class, you would likely have gotten an F and this note saying, "I think you should be a little more explicit in Step 2."

So, one solution to e-mail security is you just opportunistically turn on TLS. In other words, you try and see if you can connect to the other guy using TLS, and then you go. That helps you to prevent man-in-the-middle attacks, but it's still vulnerable to things like BGP hijacking and DNS forgery and then STARTTLS stripping, where you actually downgrade those connections. So, it doesn't work perfectly. However, it's better than nothing.

So, this is actually a graph of Gmail, which has actually done this. So, it's on their STARTTLS outbound growth, and it's actually at 90% as of October 11th. This is their inbound STARTTLS growth for people connecting to them and sending it least encrypted, even though it's not authenticated. That's at 91% at the far right-hand side of that graph.

So, better SMTP security goals besides just turning on opportunistic TLS. You have to be resistant against attacks. You must be downgrade-resistant, even on first contact. You must support a mixed environment. You must signal with your peers when to encrypt. And, you must indicate how to authenticate each of your peers.

So, SMTP is not like HTTPS. In HTTPS, you go to a website, and then you have the choice, when something goes right or wrong, of continuing. The user actually has the ability to say, "Yes, I'm going to keep going even though it fails," or, "I'm going to turn it off because I don't like their CA."

That doesn't work with mail hosts because there's nobody there to click "okay" and to keep the connection going. You can't avoid trusting them all. That's what it really comes down to. You've got to trust every CA that ever produces a certificate to be able to interact with everyone.

So, this is where the DNS-based Authentication of Named Entities, or DANE, comes in. In SMTP, the presence of a DANE TLSA record indicates a few things for you. It indicates that they do support TLS, so you must expect a STARTTLS capability, and, if you don't get it, then someone is probably trying to get you to do unencrypted when you should be doing encrypted.

You also need some parameters to do the contact. I'm not going to go into the cryptographic details of how DANE and DNSSEC work, but it works a lot like DNSSEC. Everything gets proven, not just authentic, that it hasn't been modified in flight.

Then, DANE also gives to the end SMTP servers what certificates and/or certificate authorities should have been used to sign the certificate you're about to receive over SMTP.

So, that's the high-level overview of DANE and how we got here and how it protects. It's really the only solution for truly protecting mail, even on first contact. You have to use DNSSEC to do it.

So, this is on DANE usage and operational practice. So, if you're going to go and deploy DANE for your mail servers, here's some sort of quick Getting Started tips. I have 35 slides I'm going to talk to today, or a little bit less. There is 65 slides in this slide deck. All the rest of Viktor's fantastically useful information that wouldn't fit in 30 minutes is at the end. I encourage everybody to read it because there's a lot more material. But, this is the high-level summary.

First off, to co-exist with DANE, you have to have DNSSEC-signed domains, of course. Without DANE, what happens is that, if a site doesn't have DANE, if they're not deploying DANE, like many people aren't right now, the DNSSEC Denial of Existence support comes in and says, "They don't have DANE. You don't need to assume that they are secure."

So, this actually works really well, and really, DANE is the first protocol that requires reliable Denial of Existence. The other word for Denial of Existence that you may have heard a lot of is NSEC. NSEC and NSEC3 are the capabilities that make this possible. That allows you to know whether or not to expect that STARTTLS connection, to except a secure connection and fail if you don't get it.

So, what happens when a lookup fails? Well, when protected by DNSSEC, that's when you know that you need to skip those MX hosts. You can't connect to them because you are expected to start TLS, and, if they're saying, "I don't do it," you know something's wrong and you shouldn't continue.

If all MX hosts are skipped because there's some problem – let's say your zone didn't get signed on time – remember that mail is a preferred protocol that will just queue it for a whole, so it actually probably won't hard fail unless things fail for, often, up to a week.

If you're going to adopt DANE, the hardest thing about it is actually the DNSSEC. DANE is actually not that hard itself, and coordinating TLSA records with your cert chain may look that hard, but there are a whole bunch of tools that make it easy. So, once you've actually deployed DNSSEC, actually deploying DANE on top of your mail servers is fairly trivial.

So, to do that for outbound DANE – in other words, if you have a DANE-enable Mail Transfer Agent, which is what MTA stands for, that includes Postfix, Exim, and Cloudmark – you can simply turn it on. You have to have a DANE-validating resolver close to your MTA. You have to read your MTA's documentation about that. For Postfix, for example, as long as you have a validating resolver running on the same machine, you're good to do. Then, you just enable DANE as documented. Usually, it's a few configuration options. It's not that hard.

You can do this even if your zone isn't signed. You can actually turn this on for outgoing connections even if you don't have DANE or DNSSEC enabled for incoming mail.

For the inbound path, you need a STARTTLS-capable SMTP server, which is most of them. You need DNSSEC-signed MX records. So, you have to have your zone signed. You need DNSSEC-signed TLSA records for each MX host.

If your MX host is outsourced – in other words, I own hardakers.net. I happen to run my mail server, but a lot of people don't. They'll outsource their mail to – we'll actually see a list later – to ovh.net, for example. It's one of the most popular outsource mail connections. They have to sign their end, too. So, both zones have to be signed, both you and your provider in that case.

Then, you have to always make sure that you do properly manage key and certificate rotation, which I'll talk a little bit more about in a minute.

So, there's a bunch of DANE tools that actually make this easier. You can go pull down this presentation later, of course, and look for it. Hash-slinger is actually written by Paul Wouters, who's somewhere in the room.

There's a bunch of other tools. Viktor Dukhovni has one called danecheck. In the later slides, if you have none of these installed, Viktor actually has note on how to do with just OpenSSL straight installed by itself. It's a little more complex and requires a bit more command line tools. But, the rest of it make it fairly easy just to say, "Here's my mail server. Got tell me the right DNS entries to put into my zone, and resign it," and you're done.

Next, I'm going to talk about Viktor's DANE SMTP survey. So, he has taken a whole lot of effort to collect a bunch of statistical information. He mails around once a month to the DNNSEC Deployment list a whole bunch of information, not just about DANE but also about DNSSEC and its usage according to all of the data that he collects.

So, one of the things that people have kept asking him is, "Can you please put this data on a website? Because I want to be able to look it up more regularly and not just once a month and find it in my old e-mail."

So, together, he and I have put together stats.dnnsec-tools.org. So, this is now up. I was actually working on it a little bit more thing morning. There's more to be done here. It's not quite perfect. It's not entirely automated, but it's getting closer.

Our goal is to have daily updated stats on how his survey sees DNSSEC and how it sees DANE. So, it's created, jointly in cooperation, by both of us. Keep looking at it because it's going to keep getting better. But, go look at it now if you're interested. So, that's stats.dnssec-tools.org. As I said, it's reporting both DANE and DNSSEC.

So, this is what the website currently looks like. You can see that there's a pretty graph on it that shows you the current deployment of DANE. You'll find that there's five graphs on it now, as well raw numbers if you want to look at those as well. Most of the rest of the data that I'm going to have in the following slides are actually all available there, too. So, if you ever want to see what it's like tomorrow, you can go there, rather than wait for Viktor and I to update these slides and give them again.

So, in his DANE DNSSEC survey, he monitors domains from the delegated public suffixes, and then he notifies operators of botched key and certificate rotations, which is a great public service. You'll get mail, saying, "You blew it. Here's a problem."

He used data sourced from a lot of wonderful people that he'd like to acknowledge, both ICANN's CZDS database, as well as Verisign and open access for a number of TLDs, like .se, .[eu], .fr. and .[nl]. He really wants to get more ccTLD information, so you can come talk to me and I can put you in contact with him, or you can write to him directly. His e-mail address is on the front. And then, most recently, Farsight Security has donated a lot of data to him as well.

So, some of the jumps in his graphs actually come from him pulling in a new data source that allows him to get a whole more information. But, he's covering 200 million candidate domain names, and he's looking for things like DS, DNSKEY, MX, A, Quad A, and TLSA records and then pulling all that down and analyzing them.

And, he's capturing the certificate chains of MX hosts. So, once he finds an MX host, he'll actually connect to it, pull off their certificate, and then see what matches and things like that.

So, the survey statistics. As of October 11th, when he and I talked about it, we updated these numbers, so they're fairly current. There's 8.95 million domains with DNSSAC-validated MX records. That's fantastic. That's a really big number. So, when people come around and they say, "Well, DNSSEC isn't deployed," it's true that it's not everywhere, but that's an awful large number of people that are using it.

There's 323,000 domains with DANE SMTP turned on. We'll talk more about the discrepancy about those two numbers in a bit.

There's millions of users. Comcast.net, web.de, and gmx.de all have it enabled. Comcast was actually one of the first really providers that turned on DANE support.

I think that's mostly your fault, right? You turned it –

UNIDENTIFIED MALE: No.

WES HARDAKER: No? That wasn't you?

UNIDENTIFIED MALE: [inaudible]

WES HARDAKER: Okay. So, some of the early supporters [for] Comcast are here.

So, there's 5, 538 DANE MX hosts in 3,641 zones. We'll talk more about that in a minute. And, there's 500 domains with TLSA lookups with problems.

So, when you have a growing technology, of course, you also have to look at where are problems are occurring. That occurs with any technology, whether it's security or not. But, security always adds a little bit of complexity, so you have to watch out for the failure cases.

There's 258 domains with wrong TLSA records or NOSTARTTLS. Remember that I said, if you publish a TLSA record for your mail server, connecting clients are going to require that you support TLS. So then, if your mail server doesn't do TLS, you'll actually not get the mail.

In terms of top TLDs, these are DANE domains in the thousands. So, congratulations to the Netherlands, who has three million DANE-enabled domains. Then, .com is next, with 935, and .se with 820 and so on.

In terms of reliability, breakage is usually at the parked domains. There's many lame delegations in those cases, and they're actually often just ordinary DNS outage. In other words, it's not DANE or DNSSSEC. They publish DANE or DNSSEC records, but the whole zone has other issues.

There's the Denial of Existence problem only at about 500 domains. We'll see a list in a minute of the low-breakage TLD graph and the high-breakage TLD graph. The full list is on stats.dnssec-tools.org if you want to go look at it, but the TLDs with very low DANE breakage is actually one internationalized domain, which I don't know how to read, so I'm afraid that I can't.

.br is 0.04%, which is absolutely outstanding because they also have one of the highest counts. So, even though they have a huge number of domains that exist – is there anybody here from .br?

UNIDENTIFIED MALE: Frederico.

WES HARDAKER: I'd love to talk to you later because your success is outstanding and Viktor wants to know how you did it. So, I figure – maybe I can ask you. I actually have time. Do you have a monitoring system turned on where you're able to contact your –

FREDERICO NEVES: Yeah. We basically monitor every single delegation on a daily basis, and we report on a monthly basis. So, that's the way we keep those numbers low.

WES HARDAKER: Okay. You reach out to your registered clients with mail and stuff?

FREDERICO NEVES: Yes.

WES HARDAKER: Fantastic. You nailed it, so well done.

.hk has 0.06%, but they actually have a whole lot less registered domains.

There's some high-breakage TLDs, too. Dot-bank is one of the highest. If there's anybody here from .bank, you don't need to say anything publicly, but do come to talk to me because Viktor would love to talk to you.

.nrw and .rn and .ru (Russia) is also fairly high.

Again, the full list is on the website. That's actually one that's already set to be updated daily. So, that list we'll keep up to date on a very regular basis.

The number of domains using SMTP and DANE has grown drastically since … I'm trying to see the scale on his – oh, so the years are first. So, the scale goes from 2016 to 2019 on the far right-hand side, and it's up at 325,000 domains using DANE and SMTP.

You remember how I said that most people actually outsource their domains. So, if we look at the next graph, which is the number of zones with DANE MX hosts, it's much smaller. It's 3,600. That's because a lot of those larger [3,600] are all pointing at names within this graph. So, the number of MX servers that are actually serving DANE is significantly smaller because they're being reused a lot as mail servers.

Oops. That graph was the one I was talking about. I'm sorry. So, that's the one with 3, 600.

The list of well-known DANE domains that are actually using it growing constantly. He gave this talk at Puerto Rico, and there was about half the number of domains on it. There's a couple of big ones. Of course, we should shout out to debian, freebsd, gentoo, ietf, isc, netsbsd, openssl, samba, torproject, web.de, comcast.net, and dns-oarc. A lot of these are highly – xfinity.com, for example – fantastic works that are well-used mail servers. They're not just little tiny things. They're well-used mail servers.

There's a whole bunch of almost-DANE domains. So, what Viktor labels as almost-DANE domains are the case where the zone is actually signed, but the thing they're pointing their mail servers to is not. So, this is where ovh.net comes in. Google doesn't sign there; so, people that are outsourcing their mail to Google or to ovh or one or any of these lists. The original zone is signed, but the mail server side is not. Unfortunately, that doesn't work. These are in the numbers of the left.

So, 1.4 million MX records would point to ovh.net, but ovh.net is not signed.

So, Viktor’s list of things he most wants help on are more ccTLD lists, the number of signed delegations. The more data he can get his hands on, the more he can actually use that to actually help the world and send mail when things go wrong and to update these graphs and stuff like that. That’s why, like I said, some of the jumps in the graph actually come from him pulling in new data sets and things like that – so, those big stair steps – or, sometimes, from people actually enabling a whole bunch more, and sometimes it’s from him actually being able to see a lot more.

Of course, the most help he wants is: fix things, please. When DNSSEC breaks or DNS breaks or DANE breaks, make sure that you check on things. And, check on Denial of Existence. People don’t think of Denial of Existence being as important, and, with DANE, it really is because that’s what signals that you are capable of doing STARTTLS.

Please enable DANE outbound. Again, if you don’t have a signed domain, you’re not signing your personal domain, you can still use DANE to authenticate to other servers without signing your own domain if you’re unable to do that at this point.

Then, of course, please enable DNSSEC and DANE on hosted MX servers, especially when you’re hosting thousands of signed domains, like ovh.net or Google Mail or things like that.

As I mentioned before, there are a lot of extra slides beyond this one, so, if I scroll through them, I'm not – oops. If I scroll through them, there's a lot. There's all sorts of interesting information in here with lots of other stuff. So, go look through them.

I think we're saving questions until the end of the panel? Or you want to go now?

JACQUES LATOUR: We have time for questions now.

WES HARDAKER: Okay. Does anybody have questions on that now?

[JOHN]: I don't have so much as a question as testimony. I ran into Viktor at the M3AAWG meeting two weeks ago and he shamed me into saying, "Why aren't you blah, blah, blah?" So, he has some shells scripts, which will extract the good bits from your keys and produce TLSA records. I turned it took me, like, an hour to start with those scripts and put them into my DNS management stuff, once I discovered the essential fact that, if you're using Let'sEncrypt, every time it resigns, it resigns the same key. So, you don't have to change your TLSA unless you do a deliberate key rotation, which I saved for later. So, I fussed for an hour, and now I have 90 signed mail domains and 43 signed web servers.

Knowing how Viktor tends to understate some of the difficulty, I was really surprised at how easy it was.

WES HARDAKER: Yeah. That's a very good point. Actually, let me throw up one slide. The second slide in the appendix is actually one of the biggest. It's called Rolling your TLS Keys. He gives some hints for, if you're going to change you mail keys, there's safe ways to do. He suggests two different routes for doing it. So, that's a good one to look at.

But, you're right. When DNSSEC first came out, when DANE first came out, they were not easy to able to deploy. But, these days, we wrote a tool at Parsons that you call Zone Signer and your zone file name and it's done. It generates keys if you need them. It takes care of it. It generates three keys, actually. It's much easier to use. You're right.

[JOHN]: Yeah. So long as you have any DNS automation at all, it was surprisingly easy to add.

WES HARDAKER: Thank you for that. Any other questions?

JACQUES LATOUR: Well, more an observation. So, one thing we do at .ca is we sent Viktor the list of all the domain names that are signed on a bi-weekly basis so he can run his own stats.

WES HARDAKER: Yeah, and he greatly appreciates that. And, it's not just even looking at stats. He is trying to figure out who to contact when things go wrong.

Any other questions?

JACQUES LATOUR: Alright. Thank you.

Alright. So, next up is Peter Koch from DENIC, and he's going to talk about, I guess, the deployment of DANE in Germany.

PETER KOCH: Yeah. Thank you, [Jacques]. I only have a few couple of slides and won't need all the time. So, I think we can go back to the backup slides provided by Wes and Viktor to get into some of the details. Actually, I picked [on] some of the details that I explained in more detail on the backup slides already so we can get in there.

First, if we can get the right zoom level. Maybe zoom out a bit more.

Yeah. Yeah, that should do. Okay.

So, first of all, I have an error on the title slide, which is a record, I guess. So, it should not use "use." There was a bit of slippery use of the word "use" because I can't talk about use. I don't know how many of the connections are actually used. That is, of course, a problem that Viktor also has, and Wes touched upon this. We don't really know how many mails go over these secure channels and how "important" in terms of volume the different domains are.

So, what I've done here, because we are a bit reluctant to give zone data out of our hands, basically maybe not in that in-depth way that Viktor does it, is gone throw the signed domains – wow, this shows up really great.

So, no, this is not – yeah. Now find out the difference between the left and the right black button.

For those of you were here this morning. I mentioned that we have 16.2 million domain names, and short over 100,000 are signed. So, if we look at the whole .de name space, there's only 100,000 candidate domains for DANE to start with because those are the ones which are signed. That gives part of the picture.

However, if the main mail domains were underneath, that would be important. We don't know much about those. I think Wes mentioned gmx.de and others, and they are signed. If we were able to count the number of messages, it would significantly contribute to the volume of mail that is actually transport-secured.

Now – this is a bit better.

Of these 100,000 domain names that are DNSSEC-signed, 86,000 do have least on MX record at the zone apex. So, that's, again, the next candidacy for DANE-secured mail because, I think, Viktor and Wes called it almost-DANE domains, the other ones. But, let's be a bit on the luxury side here and not only look at which of the MX records targets has a TLSA record attached to it but also look at the domain that was on the right-hand side of the [@] sign, the mail domain. The MX record should exist and signed to start with. So, that's 86,000, roughly.

Just out of curiosity here, the targets – I know Viktor loves talking about the TLSA parameters … well, the usual distribution, I'd say. There's Type 2 and Type 3. I've collapsed these believe because which Hash algorithm is used for the TLSA record isn't so much of importance. But, to look at here, the interesting one is one in the beginning. Indeed, somebody managed to get the cert into the TLSA record, actually, and then was driving EDNS0 to its extreme or to its boundaries. The rest looks pretty much benign. Of course, the [31 X] parameter we'll be getting back to in the slides explaining the details in a couple of minutes. I guess that's a usual thing.

So, the number of MX targets with TLSA is only 7%. That's not too many. I've been looking at all of the domains, so this is all .de domains. We started with 86,000 .de domains. Independent of where these MX records point, whether that is within the .de name space or outside in some other TLD, we end up at only 7% that actually do have MX targets with TLSA records. There's something to be improved, of course, and I'll come to that on the next slide after that. Which means we have 1,300 MX targets – that's the 7% of targets – and they cover 28,000 domains.

However – this is, again, about the almost-DANE domains – the next two MX targets that don't have TLSA yet would add another 28,000 domains to that. So, that sounds easy, and somebody just mentioned that this difficulty can be underestimated. Those people who were here this morning might remember that I said that much of DNSSEC deployment on the signing side is done by registrars that do full service. They operate the DNS for their clients, and then they point to their own mail infrastructure for the mail system.

So, in these cases, they operate and sign the domains of their customers, point them to their own registrar's mail server, but where the domain server resides is not signed yet.

Then, depending on the size of the registrar and the complexity of the name space and complexity of responsibilities within that company, that might be a bit difficult. So, it looks like doing one or two phone calls – have Viktor do one phone call, maybe, or one e-mail – is a bit more complicated, and we're trying to motivate registrars to, or suggest that this might be an interesting additional feature – but, there are lots of, say, non-technical influences that might inhibit the immediate deployment of these other 28,000 domains.

So, it's not all on the technical side. Sometimes it's politics, and, especially at bigger companies, responsibility boundaries that make things difficult.

So, the rest we'll only have on audio because I want to respond to some things that were already mentioned.

So, it looks quite straightforward to get the numbers higher. Again, we don't really know how many messages are flowing. Not all the domains will account for the same number of messages and users.

However, while getting DANE deployed or getting the TLSA record in the system is quite easy. Maintaining processes for monitoring key rollovers and so and so forth – [John] just mentioned that Let'sEncrypt will just sign the cert, resign the cert, over and over, and you may or may not want to roll the key. Getting these things all in place require a bit more effort. There's lots of automation underway already, but mistakes, of course, are hurting people because they are running the risk of losing mail, which is why some of them are a bit more careful.

The other thing to understand is that deploying this still, if you do your homework and install all the monitoring around this, is comparably easy. Debugging of issues turns out to be hard. Viktor has been doing a stellar job in identifying problems and helping people, especially with debugging very subtle DNSSEC issues with NSEC3 records and the like.

But, it's interesting, and, yeah, fantastic what diverse software is out there in the field and what people have managed to do with NSEC3 records – signing on the fly, white lies, all at the same moment, which, for example, in combination with aggressive negative caching, leads to interesting side effects.

Since TLSA records are not following the user pattern– www.secondleveldomain.TLDA(or QuadA) – they are of course a good debugging tool or a triggering tool when it comes to empty non-terminals and so on and so forth. We've discussed this a couple of times in these panels.

Going down in technical issues that really involve an expert, I would say that many, many kudos to Viktor for proactively reaching out here. That's very helpful.

That's all I have to add in terms of numbers. I'd be happy to cease here so that we can get to probably a bit more of technical details on your backup slides, if the panel time allows.

JACQUES LATOUR: So, we have lots of time. Any questions for Peter?

Sure, Russ?

RUSS MUNDY: Peter, I saw something in journals a couple of months ago, maybe six months ago, about some legislation that was passed in Germany relative to encrypting e-mail. Did I misread that, or is there something going on in the legal policy space that would cause this usage to grow?

PETER KOCH: So, no, you did not misread it. It's very close. So, what happened is not legislation but that the Federal Institute for Information Security issued guidelines for the secure operation of mail servers. They were targeting ISPs, main operators, and so on and so forth.

While this does not have immediate legal binding, a recommendation into that direction has certain weight. As we know and as is probably is known in these circles, the BSI, which is the three-letter acronym for that federal information security institute, is very much in favor of DNSSEC and DANE, so they included these particular protocol elements in their recommendation.

Again, that is not immediately binding, but then there is lots of IT security legislation in parallel. Many people have heard about the European NIS directive that is translated into national law. Then, there are all these clauses that would mention – I know I am probably mistreating the legal term – state-of-the-art technological measures. One may or may not judge that: "Okay. If there's a recommendation to do this, this would be considered state-of-the-art."

This remark of mine now does not constitute a legal assessment of me or my employer, but you get the message.

Thanks, Russ, for pointing that out. That's, of course, on the policy side, a very, very important addition, which is consistent in what they've been doing before, and, I guess, what's also happening, in the Netherlands, or has happened there.

RUSS MUNDY: So, have you been able to detect any changes from the technical side of things, actually, and impact from this piece of policy?

PETER KOCH: Well, not really. The reason is that, what we've been measuring here, what I've been showing a couple a minutes ago – only measuring the number of signed domains and MX targets would account for anybody – the big ISPs and everybody and their dog in their garage (or garden in that case, hopefully).

What we all, I think, would like to have would be some indication of how many mails flow over encrypted channels and, even more so, over DANE-secured TLS channels. We don't have access to that information. So, to that extent, that is hardly measurable.

What could be done, of course is, if big ISPs announce that – and there have been initiatives a couple of years ago already, where some German mail providers came together and had what I guess was called E-mail Meet in Germany or something and had negotiated encrypted channels between them, which is course kind of an in-house circle. That helped increase things a lot.

In addition to that, they also would have gone and encrypted all their IMAP sessions between them and their customers, or forced "customers," onto encrypted channels.

Again, we don't know how many mails are flowing there, but it really also gets a lot of attention by the media – well, the IT media – and interested users.

Thank you.

JACQUES LATOUR: Alright. Thank you, Peter. Questions

Oh, Jaap?

JAAP AKKERHUIS: I think I mentioned that a couple of workshops ago, but, in the Netherlands, the government has comply or explain lists, where it's actually stated, for procurement of services, what minimum things you should do. DANE is on the comply or explain list. So, basically, if you want to have a government contract, you should offer DANE automatically, or have a serious reason why you don't. So, there's quite some strife about it. Only European countries have seen to take off with this idea. I[inaudible] because that's always filling the box and calling your lawyer exercise. But, they needed to get stuff done.

[JACUQES LATOUR]: Wes?

WES HARDAKER: One thing I think we haven't mentioned today, or I failed to, is that DANE is actually a generic mechanism for securing TLS connections. We are talking about SMTP because that's really where it has taken off.

Have you done any measurements of web service or things like that? The web browsers aren't really using it, so it doesn't make a whole lot of sense. But, is there any other TLS-based protocols you've looked at for –

PETER KOCH: Quite frankly, no. We didn't do that for the obvious reason that there's little to expect from web browsers to follow up. So, some people are doing this just out of curiosity because they like DANE. Maybe they're even surprised that web browsers would not adhere to what information is in the DNS. But, we didn't look into it more deeply.

WARREN KUMARI: Wow, that's loud. So, I should point out that there is a new draft in the IEFT – or, a fairly far along draft in the IETF, actually – which takes all of the DNSSEC information and stuffs it all in a certificate, and then all of the DANE information – I guess I should reword that – and stuffs it all in the certificate. That allows a web server to provide that to a browser. That means that the browsers have to go off and do the DANE work themselves. They don't have to do additional DNS lookups.

So, yeah, as you were saying, this can be used for more stuff, and the TLS extension certificate thing may actually get browsers to do this. Or, maybe not.

WES HARDAKER: Do you know if there's been any commitment by browsers to want to do it?

WARREN KUMARI: Umm …

WES HARDAKER: Okay.

WARREN KUMARI: There may or not be. I cannot comment on that.

JACQUES LATOUR: So, what is it called – the TLS thing?

WARREN KUMARI: So, one of the stated reasons that various browsers have not currently done DANE is that it requires doing a bunch of additional DNS lookups. Some of these will fail because some DNSSEC lookups always fail. Some of them add additional latency, etc.

What this draft largely does is it takes all of the DNSSEC information and stuffs it in an extension in the certificate. That way, when you get the certificate, it has all of the DANE information already, and the browser or whatever other application, doesn't need to go off and do the DNS lookups themselves.

Actually, now that I think of it, Paul Wouters is sitting next to me, who actually is involved in this.

PAUL WOUTERS: Hi. So, you're mostly right. There was an old way of doing this where they would stuff into the [inaudible] certificate. But, that didn't go anywhere for other reasons. Now, it is being stuffed in a TLS extension, so it's not a certificate. But, it's a little bit important because it's actually much more separate from the web PKI than that proposal is.

Some browsers wanted to implement it, but there is a huge discussion happening now in the TLS Working Group about moving this document forward and why it's been taking nine months of no new draft.

So, if you're interested in that, there's 1,500 e-mail message on our TLS Working Group list. But, if you walk up to me, I can give you some summaries of that.

But, people are interested. Work is happening. But, it's a little bit frozen right now. Hopefully, in two weeks it'll move forward again.

JACQUES LATOUR: Peter?

PETER KOCH: So, I'm not turning this around and asking the audience from the panel – and this is only partly DANE-related – but, if we take e-mail transport security as the topic, if I recall correctly, there's also work on the way in the IETF on strict transport security for SMTP. If memory serves me correctly, an additional project that is addressed is signaling the transport in the message or to the message receiver or something, which would help at least the occasional user to find out after the fact if the mail was transported securely or encrypted, I should say, or not.

Maybe somebody could elaborate on that. That would help a bit to make users understand what effort is actually about and whether or not they are a part of this and maybe should talk to their ISP or in a similar direction.

WES HARDAKER: So, I'll speak a little bit to that, which is that some of those are the equivalent of Trust on First Use. So, there's some movement to do pinning. That really doesn't help you a huge amount, certainly on the first use. But, the thing is is that, if somebody can insert MX records, they can point your mail to anywhere. So, if they could mange to get cache poisoning or something else, if you don't have DNSSEC turned on, they can change your MX records. To at least the person trying to send you mail, it's going to go somewhere else first.

Then, they could deliver to you, anyways. So, you'll actually get a copy, and you just won't know that it was seen by a man in the middle. It's one of the worst attacks.

So, when Viktor and I worked on the DANE SMTP draft, you have to start in DNS. So, by having the DANE record in DNS, it prevents all of the man in the middle of the stuff that the rest of those solutions don't quite handle. And, it works on first use. So, it's really the best, technologically speaking.

JACQUES LATOUR: Alright. So, I guess one observation is that we're in a chicken and egg situation. So, the browsers are not doing certificate pinning with DANE because there's no DANE record available to do certificate pinning to test with.

So, if we add more records available, and then it shows that it's feasible and it works – if we sign more domain names and have TLSA for the website available, it would make … even though the browsers don't do it, we should turn it on? Right now, it's not turned on.

WES HARDAKER: You said sort of two different things. So, I think what you mean is that, if we publish TLSA records for web browsers, maybe it would indicate … and somebody said [they wanted] to be in the past, that they're watching, and if there's enough deployment of it, they'll turn it on. But, I don't think that's the truth, unfortunately.

But, there is – I think we did DANE in Bloodhound, didn't we?

UNIDENTIFIED MALE: Yes.

WES HARDAKER: So, there is one web browser that supports DANE.

WARREN KUMARI: If enough people publish DANE records, there's a reasonable chance of the TLSA extension thing, or there's some chance of the TLSA record thing, browsers might be willing to do. I think it's very unlikely that browsers would ever be willing to do the actual DNSSEC DANE stuff. It simply adds too much latency, and some number of users aren't going to be able to get DNSSEC-signed records at all because of middleboxes.

So, [that] the TLSA record or DANE record stuff/certificate extension might fly the actual DNSSEC lookup for DANE records is very unlikely, too.

WES HARDAKER: So, there was a discussion at one point about, well, you show them the page, you start loading it immediately, and when DANE fails eventually, then you take the page away from their eyes.

WARREN KUMARI: That was my idea. I pushed it and the browser people and the security people that that's a bad idea. And, they're right. By the time that you've done that, you've already shipped all of your cookies off to the bad site, in which case it doesn't really matter because they can emulate you anyway. So, they were right. I was wrong.

JACQUES LATOUR: I like that idea, but, Alright. Any other questions for them?

So, the next presentation is from Wes, again, and it's Speed Up Your Resolver DNS With LocalRoot.

WES HARDAKER: Once the slides are here, we'll start. So, I guess my first plea would be that everybody should work on presentations for the DNSSEC Workshop so you don't have to listen to me three times in one day, although this last one was sort of a last-minute change and I'm really just a filler. But, thanks for sticking around.

So, I think it was about a year ago when I talked a project of mine called LocalRoot. The whole purpose of LocalRoot we'll get into in a minute. This is a really an update on new things that have occurred, new pieces of it, that are helpful. I'm not going to need 20 minutes.

So, this is a screenshot – can you scroll left just a little bit? Because we're going to miss all the beginnings of the titles, otherwise.

Thank you. So, this is a screenshot of the website. We'll get to it in a minute. It's localroot.isi.edu, if you want to go look at it right now.

The purpose of LocalRoot is to do a couple of things. It's supposed to do pre-caching of DNS records. Originally, it was only root zone data. One of my updates today is that there are actually three zones there. I would like to do more. I'll talk about that in a bit, too.

I really wanted to make it super-easy to deploy. So, if you know how to run a name server, this should be trivial because I'll actually spit out config that you can just copy. You don't have to do anything other than copy and paste.

It supports 7706. So, it's RFC7706, or, in the slides up there, RFC770U – no, it still left off the six. Wow.

UNIDENTIFIED MALE: It's a feature?

WES HARDAKER: It's a feature, yeah. So, why would you want to use LocalRoot? So, I should mention what it is. I originally did these as just an update, but it has been a year. So, what LocalRoot does is it allows you to have a copy of the entire root zone pre-cached in your resolver. That's fairly easy to do today anyway, but this actually gives you notification and it gives you TSIG ability so you actually make sure that the connection between you and the server you're pulling it from is secure.

By doing that, there's a number of benefits. In particular, you get local service meetups. So, by having the root zone in your local resolver, you never need to ask for it. If you don't ask questions, if you already know the answer, then it speeds up your resolution quite a bit. And, it's true for pretty much all zone records, whether it's A, QuadA, DANE, or DS records or whatever.

And, it prevents outage. If you have a copy of the root zone, and your transport to your favorite root zone provider was done, you would still have a local copy.

So, the root zone is unlikely because there's never been an outage of the root zone. But, there has been outages of lots of other stuff. So, one of my eventual goals is to support other zones as well. So, there's three in there now, but eventually there will be more.

And, it's privacy-protecting. I have done a presentation in the past at both the last NDSS – and I may have given it at ICANN;I don't remember – where I basically showed that, even if you are doing query minimization and using TLS, just by asking questions, you're releasing data. If you just go ask for the NS records for .ru, I know you're talking to Russia. I'm not picking on Russia in particular. But, if you're talking to .bank, I know that you're doing banking. I know the time of day that you're doing banking. There's timing information.

So, the best way to protect your privacy is: don't ask questions. If you have all the answer pre-cached, it's significantly more private.

So, RFC7706 I'll give a quick summary of. But, one of the sentence of it says – and the author is over there, by the way – "The primary goals of RFC7706 are to provide faster negative responses to stub resolver queries that contain junk queries and to prevent queries and responses from being visible on the network."

If you remember, I showed you a graph earlier with a whole orange line full of junk queries. The root service gets tons and tons and tons of random junk. So does .com. If you query something for .com that doesn't exist, they get a ton of junk stuff as well.

The reality is that most queries to the root don't exist, as I showed in the graph previously, and that negative answers are cached for a shorter time. So, if you go ask the same questions a few minutes later, you won't keep it in your cache because it wasn't real data.

7706 can really help with this if you're on a slow link, if you're using a dial-up modem, if you're in a distant part of the world, and even on my – I won't admit what network speed I have at home, but it's not stellar. Enabling root on my home box has actually helped some.

DNS traffic is observable to on-path third parties. So, that's the other thing that I mentioned about privacy. That is one of the other goals of 7706: privacy leakage of possibly sensitive domains. Again, 7706 means that no questions actually traverse the network in the first place.

So, this is where I get to the update portion. A bit of news. There's four slide of news. Most importantly, recently I had to change the name of the TSIG keys that you link to my server with.

Unfortunately, TSIG keys have to be A) both unique and B) can't change. Well, I just changed them, so, if you're using LocalRoot, I do need you to go change your TSIG key. You'll go get a new copy of the config. We'll talk about that in a minute. This is the last time I'll do it. I've guaranteed that this is going to be unique this time, so I won't need to do it again.

It's now possible to actually delete unused servers, unused TSIG keys. That was actually the number one request people gave me. "How can I delete stuff?" "Well, I haven't implemented it yet." Now I have.

There's now new e-mail account notification preferences. So, when you sign up, you can actually click on boxes saying, "I want to receive updates," and, "I want to receive bits of information." I haven't actually sent out anything yet, but I will shortly. And, there's a whole ton of user interface improvements to make it easier to use.

By the way, the one thing I should mention is that the reason this is being presented in the DNSSEC Workshop is that, before DNSSEC existed, this whole project couldn't have happened. You couldn't have trusted the data. This really enables you to be able to data sources from anything that's signed.

We also support two new zones. .arpa and root-servers.net, which are the other two zones that are also served by the root servers, are in there now.

There's a lot server screen improvements. One of the other things that's really helpful from a debugging point of view on your side is that the last transfer time is shown. So, the smallest font on the bottom left -- this is a screenshot showing that that server is both active and it's enabled. There's a git config button on the right, which we're going to go into next. But, on the left is a time stamp that shows you the last time you pulled the zone so you know that your server is actually pulling fresh copies.

There's a much better configuration generation screen. It's now dynamic. You can put in all sorts of information. You can put things in like, "Do you want the full configuration for BIND, or do you want a partial config that you can just cut and paste because you already have a BIND config you need to insert stuff into?"

It also automatically includes your private address spaces. If you are an ISP resolver serving private address spaces behind you, like 10.whatever or 192.whatever, there's check boxes, so it'll automatically put that configuration in so that you will enable service of it.

Then, of course, one of the other things that I realized is that I was defaulting for a while to where you're going to store your zone file storage, where your slaves were going to transfer zone files to. And, that didn't exist on debian's, so I fixed that so you can actually specify a "where" to store the files that it has to download.

That's just sort of the beginning of stuff that I really want to do. I'm probably going to write a small grant application soon to get some of my students and other people to help me do a whole bunch more interesting stuff, like improved infrastructure. Right now, there's really one DNS server that will transfer the zone to you. If it's not up, you'll actually get it from the rest of the root servers, too. But, the TSIG version will only work with mine.

It doesn't support IPv6 yet. I know people want IPv6. That's hopefully coming soon. And, I want to be able to send e-mail notifications. When you fail to do an AXFR recently, you'll get a note saying, "Did you turn it off? You ought to come back and disable it, or maybe, "Something is going wrong." The same thing is if I see failed TSIG connections because something was wrong with the key. I want to be able to send e-mail notes on that.

I want to support Unbound and other resolver configurations. That's actually one I don't need to go find a funding grant for. It's actually quite easy. I just haven't done it.

Other sources of DNS data. This is actually my largest desire and biggest project that takes more work, which is to be able to pull in a selectable list of, if you use a TLD that we have data for, let's go pull it in. You probably don't want big ones. I can't get .com for you anyway, but you couldn't pull it in. Your machine probably doesn't have the memory for it. .com servers are pretty impressive.

But, there's a whole lot of TLDs which are actually quite small, and wouldn't it be great if you never had to ask them for information because you already had it?

So, there's a bunch of places I know I can go get a lot of other important domains, and I've talked to some vendors that want to explore this territory with me as well. So, that's coming at some point, and I'm really looking forward to that particular piece of the most.

Then, of course, I would actually like to publish all of this code. It's all written out in my local git repository. There's no reason I shouldn't push it to GitHub. I just need to do that.

Any questions? Again, if you want to try it out, go to localroot.isi.edu. If you run a resolver, it's quite easy to cut and paste name decodes straight of it. You just create servers, create TSIG keys, and click "Get con fig," and it'll dump a bunch of stuff at you. You can just cut and paste.

Any questions?

JACQUES LATOUR: Jaap, first one?

JAAP AKKERHUIS: Now I suddenly remember the talk I was supposed to give here but didn't do because I forgot. A mechanism like this is actually standard in Unbound and is in production and is not the default switched on at the moment. But, it's [inaudible]

So, we are competing there. But, it's slightly different than the way of doing stuff.

But, the other thing is, what about the [notice]? It's not that you're afraid that all the root servers are [inaudible] for a while, but what it really solves is the problem of, if you are on a little island and the only connections to the last of the world is via satellite and something happens – a tornado or whatever – then, suddenly, you can even not resolver locally any more when the root information is expired. But, by having it served from the local cache, it actually helps you at least get the local communication call less depending on whether or not to do [inaudible]. That's actually the big win for this type of system.

WES HARDAKER: Yeah. That's a very important point. That was one of the other primary purposes of 7706, especially the latency on satellite links and stuff. I actually did some research for a while that showed satellite links were one-third lost. It's quite noisy.

We're not competing so much because – you're right. I know about that. You're talking about the ICANN-sponsored project. They also did it for BIND. So, actually, both of those name servers had the ability to pull the root zone specifically.

As I said, I'm hoping to do more and to do more enabling. Then, my mechanism also gives you notifications, so you'll name servers will actually get a notification to go get root zone x or to go get root zone now or to go get the rest of the TLDs or whatever else I manage to pull in eventually. So, you'll get that, which doesn't come with the "I'm going to go fetch on a regular basis" type of approach.

ABDALMONEM GALILA: Do you image how it works if I have my own local root zone inside my resolver and I would like to sign with DNSSEC? Where's the chain of trust at this point?

WES HARDAKER: So, you are pulling the real IANA root zone with their keys and their trust. The DNSSEC still works the same. You don't put it in your zone file. It ends up being a local copy. It's like as if you had already queried the root and it already remembered all the answers. There is no security difference, and it's the same set of keys. You're not resigning.

ABDALMONEM GALILA: It's the same trust anchor?

WES HARDAKER: Same trust anchor, yeah. There's no difference, security-wise.

ABDALMONEM GALILA: I think this is a good idea for blockchain, to use DNS as a service for blockchain.

WES HARDAKER: That's a whole different ball wax. We can talk offline about it. But, this is unrelated to blockchain, by the way.

WARREN KUMARI: So, there's some work underway in the IETF – I'm not sure if I can see Duane; oh, hey – to make it so that all of the information in the root zone is going to be signed, including glue and all of that sort of stuff.

Once that happens, do you think you'll still need the TSIG stuff for transfers, or then can it just be, "Just leave it AXFR?" or –

WES HARDAKER: It definitely will matter less. TSIG feels better, but you're right that it shouldn't matter as much. With TSIG, you likely know you're talking to the right server. Otherwise, if I was in the middle, I could give you something too old. With TSIG, you know that you're actually getting the most updated zone because you can trust the person on the other side of the connection. So, I'll still keep it on.

DUANE WESSELS: I was going to ask a similarly loaded question. How can someone getting a zone file from you know that they're getting the right zone? How do they know that you're not fiddling with it?

WES HARDAKER: Well, because of DNSSEC. What I said sort of halfway through the presentation is that this really wasn't possible because you shouldn't trust my server to give you the right information. You should validate with DNSSEC to prove you have the root zone signed by the ICANN KSK.

DUANE WESSELS: Well, I was trying to get you to put in a plug for our Internet draft that we're co-authoring, the message digest draft.

WES HARDAKER: Yeah. Well, but even then, I could give you that digest, right? Well, I guess no. It is signed by the KSK. So, no, that would work as well, although there's still the timeliness issue that I pointed out before.

WARREN KUMARI: So, yes, somebody could give an older zone, but it should be fairly easy to notice that because the root zone gets cut at least once a day and the SOA assigned. So, you could at least tell. And, this has just occurred to me, so I'm making it up on the fly. So, I could be wrong. And, I'm going to keep talking so [inaudible]

WES HARDAKER: So, in order to tell, you would have to go check somewhere else to make sure that … anyway.

WARREN KUMARI: If the SOA is more than one day old, I know it’s old.

WES HARDAKER: So, to answer both of you, when that gets done, yes, LocalRoot is one of the places. And, I told people that I want to turn on the results of that draft for this as well.

BRETT CARR: Brett from Nominet. I’ve just signed up, so the one comment I have – I may missed it in your presentation – is that I really like the fact that you pull the zone from you but the backup is the root servers that support AXFR currently. That’s really good.

WES HARDAKER: Yeah. And, I should thank them all because I reached out all of them and said that a lot of the root servers support doing AXFR. All of the ones listed there specifically gave me permission. I didn’t assume that. So, everybody there, actually. So, yeah. It’s really a failsafe system. There’s no way that you’re not going to get a zone transfer from somebody.

UNIDENTIFIED MALE: Thanks for doing that work, Wes. The RFC is already a couple of years old, and I think it’s important to get some operational experience with this.

Also, the aspect of scaling would be interest. At this point in time, it’s interesting that we see DNS being more treated like content, as in what built here is kind of CDN more or less, just for the DNS and not necessarily centrally controlled, which, yeah, suggests two questions.

One is, because you suggested also to extend this beyond the root to TLDs, for example, is, how would you monitor the quality? I'm not so much concerned – well, I am concerned, would be concerned, about the authenticity and the integrity of the data, but the quality of the delivery, where a couple of TLDs would have SLAs in place, or similar things. That's one aspect.

The other is we would shift the requirements and also the service levels from the answering of DNS packets to provisioning the zone. Would that mean we need completely different infrastructure on top of that? Because I'm sure that the University of Southern California is not going to be the only place forever to actually do this. And, we do have root servers. I'm not really aware, but there might be work underway in RSSAC to deal with the sourcing side of this.

Could you elaborate on that?

WES HARDAKER: Well, I can some. This is not just an experiment because it's actually infrastructure, but it's to get us to start playing in that space that you just mentioned. The root zone is the thing that people have talked about the most. It's like, how can we give everybody that data in other mechanisms so that they already have it?

My push in this was to take it even server. Okay, it's not just the root zone. With all of DNS, you already have what we were just talking about.

This is really just kind of pre-caching. So, you're right. I am giving you an entire, say, set of end zones ahead of time, and it does change the service expectations.

One of the things I'd like to do not on that list because it's further down the list is to have a [Daemon] test suite that could actually make sure that your resolver is actually behaving properly.

But, we should have that already. We should already have stuff that makes sure your resolver is working properly. Right now, it fails, and that's when you go check things.

This is no different, data-wise. You either have the data or you don't, in your cache. The question of whether the cache is a slave zone, in BIND speak, or that you're pulling it every time is just a question of where it is.

JACQUES LATOUR: So, two. I'm not sure which one is first.

UNIDENTIFIED MALE: Hola. Why does this not promote LocalRoot for our Internet service providers?

WES HARDAKER: Why not promote it for …

UNIDENTIFIED MALE: LocalRoot.

WES HARDAKER: So, I am trying to promote it for Internet service providers, for ISPs, absolutely. Play with it a while before you turn it on to a large number of people to make sure that you have it working. But, that's sort of my hope. Right now, it's probably being used more in small places than in big ones, but there's no reason an ISP couldn't turn it on so that an entire region has this enabled. There's no reason it can't be used that way.

JACQUES LATOUR: Daniel?

DANIEL: So, I don't get fully much where the source, actually, of the … do you do [inaudible] for requests.

WES HARDAKER: Yeah.

DANIEL: Okay.

UNIDENTIFIED MALE: [inaudible]

WES HARDAKER: Right. So, where does the data come from? So, as the gentleman over here noticed, there's actually a list of servers that the LocalRoot config gives you, including a whole bunch of root servers but that offer AXFRs anyway, like B. L does. F does. There's a bunch of them. So, you can pull it directly from them, but there's a LocalRoot DNS server that you ill also do a TSIG-secured transfer with.

DANIEL: Okay. So, to extend that to .com, for example, maybe you can have a sub-zone of .com, which only has to sign [our set]

WES HARDAKER: So, there's no reason I couldn't do sub-zones within, not a partial piece of .com, but I could do example.com as an example, or, .ietf.com or something, that would let us do it.

But, you can't do all of .com.

DANIEL: No. Not all of .com, but only those that have signed that.

WES HARDAKER: So, one of the other things that I thought of would be a tool that would do sort of some artificial analysis, not to serve it the way that LocalRoot is having you do it, but to pre-cache. It's like you go, you run a script out of cron every day because you just love IETF.org, so you go pull it once an hour to make sure it's always in your cache. You can do that. You'd need machine learning and artificial intelligence to figure out what the ones you use the most one.

DANIEL: No, but that could encourage the people to sign their domain because they're going to be pre-cached.

WARREN KUMARI: So, actually, something related. Joey Abley and I and a few other people – possibly you, actually, now that I think of it – have been talking about building a service where people could sign up, kind of a DNS exchange, where people could provide their zone files to a central location and that would become a central transfer point and anybody else could slave them if they wanted. We were planning on tying that in with a cache flush option as well so that people could report or request that caches flush a specific zone out of their cache. So, that's somewhat related.

WES HARDAKER: Yeah. I would say I have 90% of that infrastructure done for you.

WARREN KUMARI: Yeah. We'll chat more.

JACQUES LATOUR: Are we good on time? Or, one more?

UNIDENTIFIED MALE: [inaudible]

[JACUES LATOUR]: Alright. Yes?

ABDALMONEM GALILA: This means – correct me if I am wrong – that, each time ICANN tries to resign the zone, [ICB] should have a new version from the root server. The cache time at the resolver will cache all the data, so it should have another time to [inaudible]. This means that my client will be dropped. That’s my first point.

The second one is that maybe I wouldn’t trust my resolver and I would have my copy. So, every time I should get another [inaudible]?

WES HARDAKER: So, to speak to your first point, you get a new copy of the data immediately, so that’s one of the advantages. It doesn’t matter so much for the root zone because it really doesn’t change that much. But, if I can get to the point of other zones that might get more frequent updates, you might care more.

For the root zone, very rare changes go into the root zone. I’m sure some people can talk about it. When the new gTLDs come out, you would get a new one right away, but nobody is going to look at a new one right away. New gTLDs take a long time to actually get up and running and get clients underneath them and stuff like that.

UNIDENTIFIED MALE: [inaudible]

WES HARDAKER: So, the signatures you'll get right away, too, yes. All of that, the entire zone file, you'll get, including the signatures, right away. So, you'll always have the freshest copy within seconds, probably, of the …

That's actually one of the other things that I've had on previous slides: that it enables researchers to actually do some research on, say, zone update frequencies and stuff like that because they don't have to pull the root zone constantly to figure out when it gets updated. You'll actually get a DNS notification the instant that the root zone is updated. It propagates a little bit. It goes through a few systems to get to it. But, it should be on the order of seconds after the root zone is updated when you'll get a notification. So, there's interesting research that people might be able to do there because it's not on a synchronized clock every single point of the day. The point that it gets updated is not consistent.

JACQUES LATOUR: So, how do you do the roll on LocalRoot?

WES HARDAKER: So, one of the interesting things that RSSAC is actually considering too is that, with things – so, in the rest of ICANN, this technology is being referred to as a hyperlocal root. That's the key words that the Board in particular are throwing around, and other people have, too.

When you have a local copy of the root zone, none of your data will get into DITL. None of your requests will go into DITL. Some people see that as a concern because we'll know longer have the research visibility into the data sets that DITL provides. If the entire world was running LocalRoot, DITL would be pretty blank. It was just be the AXFRs, pretty much. So, some people consider that a problem. Some people consider it a benefit, wherever you fall on that or not.

JACQUES LATOUR: Well, I guess the name of the game is having the content as close as possible to the eyeball, so eventually we'll have .ca inside Firefox or something.

WARREN KUMARI: Is that a commitment?

JACQUES LATOUR: No, but eventually.

WES HARDAKER: So, I actually did talk with – I don't think it was you – somebody else from .ca that was interested in using the same thing. This is one of the reasons I want to release the code: so that they could get their local ISPs to have .ca, even though it's large, all pre-cached. So, when .ca gets DDoS, nobody cares.

JACQUES LATOUR: Yeah, that was me.

WES HARDAKER: It was you. Good. We should still work on that.

AHMAD ALSADEH: Hello. This is Ahmad Alsadeh, a Fellow. I believe the root instance would solve the problem of a local root. Instead of having local root, it can host one of the root servers near to my infrastructure or inside my infrastructure. If you agree with me, that's good. If you don't, don't you think that your project would affect the number of instances of the root servers?

WES HARDAKER: Good question. So, two things. One, yes, you're right. If you had an instance from one of the root operators right next to your box, then this probably isn't as needed if that's what you're trying to do.

However, we will never get to the point where there is enough physical hardware instances that could be shipped to every ISP on the planet. It's just not a conceivable project. So, this allows ISPs to use existing hardware – no new purchases, no new whatever. You are just installing a copy of it locally.

AHMAD ALSADEH: What is the box that does this routing to this Anycast [inaudible]? It should go outside for the nearest Anycast point of the root server and then go to the local root server again?

WES HARDAKER: Actually, that's an interesting point. If you had an instance near you, name server software doesn't always use that same box, whereas, actually, LocalRoot, because you're actually running it right there, will never ask outside the box. I've actually proven that. If you would have an instance nearby, you would ask that one most of the time, but occasionally you'd send stuff to other ones, too.

RUSS MUNDY: I've got one question for Wes. When we get better definition of what kind of statistics and data are needed to be collected from the root server to get – whether it's DITL data or something similar – is that somewhere on your roadmap, to possibly incorporate some of that or a subset of that so the information is available as more instances of 7706 show up?

WES HARDAKER: Well, that's a really interesting question. By the way, I had no idea this topic was going to generate so many questions. Thank you all.

So, there's no reason – the thing is, with DNS-OARC, anybody can contribute DITL data. There's no reason that local root operators couldn't contribute DITL data in theory. And, there's certainly no reason that I couldn't produce the equivalent of RSSAC002 stats as well.

So, it's not on my roadmap. I should add it. I will say that I definitely need to find a funding source to implement something that big. But, no Absolutely there's no reason that couldn't be done.

WARREN KUMARI: So, would you really trust a whole bunch of DITL data coming from, for example, me posting from my local root instance?

WES HARDAKER: Wouldn't that be an interesting research project in itself; to figure out whether Warren was trustable or not?

WARREN KUMARI: Would you want some funding for your thing to look?

JACQUES LATOUR: Any other questions?

Alright. Thank you.

WES HARDAKER: Thank you, everybody. Remember, write presentations so that you don't have to listen to me.

JACQUES LATOUR: Or me.

RUSS MUNDY: Well, thanks to our DANE panel and all of the interaction that we had and questions and back and forth. In fact, that is one of the reason why we host these sessions: to not only hear what's going on in the community by people that are out there doing work that are able and have the chance to get up and talk about what they're doing, but to get feedback not only on that work but on new ideas and capabilities that come from the community as a whole.

So, I think that this was a really excellent session. I don't know. We might have to go back and check the timing on that one chart from Viktor to see if the DANE usage right after he presented at the workshop last time.

Do you happen to know, Wes?

WES HARDAKER: I don't know, but I do have a piece of information for you. Since I started my thing, one person has successfully deployed LocalRoot and has it up and running and it's verified and working. Was that you?

UNIDENTIFIED MALE: No.

WES HARDAKER: Not you. Okay. Somebody in this room just did.

RUSS MUNDY: Fantastic. Okay. So, thank you for all of the presentations today to all our presenters.

Is it stuck?

UNIDENTIFIED FEMALE: This one crashed. This one hasn't.

RUSS MUNDY: Oh, okay. So, we have an interesting situation that shows up now at the end, and that is that our Adobe Connect room machine is working … I think.

UNIDENTIFIED FEMALE: Adobe is working. The presentation is not.

RUSS MUNDY: And, the presentation machine is not. So, the last element of our discussion today really is a chance for people in the room to think about what they heard today, think about what they might be interested in doing, and also to give folks a chance to think about what you might want to talk about next time because we do have these workshops at each session, and we fully expect to have another one at ICANN 63. So, we really do want to have folks leave here thinking about what not only they want to hear about but what they might want to talk about with folks.

So, thanks, Wes, for the earlier plug for more presentations from folk and from additional folks. We will be sending out – probably it'll be around the first of the year, or maybe a little before the first of the year – the request for participation. So, keep your eyes out for it. We send it to multiple mailing lists. Be thinking about what you'd like to talk about at the next meeting.

So, in terms of just what we kind of covered today and people in the room, everybody sort of has their own spot that they sit at in the DNS structure of things because it's a big and diverse structure. Think about, at a minimum, if you have any impact on a zone, whether you are a holder of 1 or 100 or 1,000, signing these zones. Get them published as signed zones. Work with your registrar if you're working through registrars.

We've reached a point now successfully with DNSSEC and the registrar community that, if your registrar you've been using does not do DNSSEC, you now have places, additional registrars, that you can sometimes use that are doing DNSSEC.

So, you can move names from one registrar to another. I have to say, personally, I've done that a time or two just because the registrar that I had them with before was not doing DNSSEC. So, this is something that can be done.

Turn in your statistics. Work with people that are collecting data, doing research projects.

Next one. So, we really are wanting to get as many people involved as possible. Enterprises – we have a number of enterprises now. Not a huge number, but there are a number of enterprises that do have DNS signatures over their published zone. I am glad to say that parsons.com, my employer, is one of those, and they have operated as an enterprise-signed zone for quite a while.

Again, [work] with the operators of not just the registrars but the name service operators because, a lot of times, it's [outserviced]. We've heard about a lot of places that are providing these services today.

So, get engaged with ways that you can make use of and start asking for more people to do more DNS things. So, if you're a service provider, an ISP, we've had great feedback today from people at our ISPs[and] want to do more. I encourage people to do that. This group, this community of DNSSEC has been excellent over the last, well, I guess 15 or 20 years I've been [diggling] with it – about sharing information. People are more than happy and willing to share information.

So, stay in touch with each other. If you need help, reach out in multiple ways, and people will do that.

So, I think that, with that, we're getting close up to the end time – yes, [Matt], go ahead.

[MATT]: Just one presentation about presentations. Will all presentation be uploaded to the –

UNIDENTIFIED MALE: They already are.

[MATT]: They are? Because I already loaded and I couldn't see Wes's presentation there.

UNIDENTIFIED MALE: Each section has its own presentation.

[MATT]: Okay.

RUSS MUNDY: So, the DNSSEC Workshop, because of the structure of the ICANN meeting, is in three parts. So, Part 1, Part 2, and Part 3. [Patrik] gave us some slides, so his might not be up yet.

UNIDENTIFIED FEMALE: I will.

RUSS MUNDY: Yeah. We'll get those up later – ah, they're there.

Okay. Anybody have any other comments or questions?

Yes, Jacques?

JACQUES LATOUR: Is it worthwhile asking if there are topics right now that people would want to [like] next time? So, while we're in it.

RUSS MUNDY: Ah, good. Yes. Does anybody want to raise a specific topic for the Program Committee to consider for next time?

UNIDENTIFIED MALE: DNS using blockchain.

RUSS MUNDY: We will include that in our call for participation. So, yeah. Be thinking about how we can do that. Okay. Great.

Another one from over here? Sure.

UNIDENTIFIED MALE: A progress of the browser in the local and other discussion in this meeting today.

WES HARDAKER: So, the web browsing world and its interaction with DNSSEC and DANE.

RUSS MUNDY: Great. Okay. Any other suggestions?

Yes?

JACQUES LATOUR: I want to ask a question. Why is there no DNSSEC Quiz this year?

RUSS MUNDY: The reality is that we had so much material this time that we elected to not include the quiz. We just had enough input that chopped the quiz.

JACQUES LATOUR: Okay. If we find a volunteer to do the next quiz, that'd be awesome.

WES HARDAKER: I would have spoken for the fourth time today to do the quiz. In fact, I already have one question that I've been saving for the next time I do it.

JACQUES LATOUR: We'll have a quiz next time. We'll have one.

RUSS MUNDY: Okay. So, once again, thanks again to our lunch sponsors: Afilias, .ca, and SIDN. And, thanks to all the presenters. I look forward to seeing everybody at the next workshop. Watch for the call for participation. It'll be coming your way. Thank you.

[END OF TRANSCRIPTION]