LOS ANGELES – DNSSEC Workshop
Wednesday, October 15, 2014 – 08:30 to 14:45
ICANN – Los Angeles, USA

UNIDENTIFIED FEMALE: DNSSEC Workshop – 15 October, 2014.

JULIE HEDLUND: Welcome to the DNSSEC Workshop. I'd ask you to take your seats. Please take a seat. You're welcome to sit around the table. And in fact, we encourage you to sit around the table, except for the spots at the front of the table that are listed as for speakers. If you are a speaker, then please do come to the front of the table.

My name is Julie Hedlund and I'll be supporting the workshop today. So we'll start momentarily with Dan York, a presentation from Dan. Let's give it a few more minutes while you all get settled, and then we'll get going. Thanks, everyone.

Good morning, everyone. Again, this is the DNSSEC Workshop on 15 October 2014 at ICANN 51. My name is Julie Hedlund. I'm with ICANN staff and I'll be helping to support this event today. I had also introduced Kathy Schnitt to my left. She's also an ICANN staff person and I'm thrilled to be able to have some assistance from staff – yay! Not just me! So anyway, Kathy will be helping out as well.

So just one housekeeping thing. There will be lunch and lunch will be right outside this room. There will be stand-up tables. There's not going to be enough room for tables for everyone to be seated at once. We have a full house. You'll have to take turns. There are tickets. You're not

really going to necessarily need the ticket to just get out of this room, because obviously the lunch people won't be able to tell where you came from. But if you were to leave the workshop – God forbid – and go do something else and then come back for lunch, take the thing that looks like a program on one side (which it is), and if you turn it over, there's a ticket on the other. I love Jeff's expressions. That's great. Thank you, Jeff. So please do retain your ticket if you do want to come back for lunch.

So that is it for me. We're starting out with Dan York. He's going to talk to us about some very interesting things. So take it over, Dan.

DAN YORK: Indeed. Well, good to see you all here this morning. We have to say as well our normal thank you to Julie for all the amazing work she's done to help us get this workshop together, so I'd like to start out with a round of applause for Julie. [applause]

As a member of the Program Committee, I can tell you there's an awful lot that happens to bring you this six-and-a-half hours that we have today, or six hours and 45 minutes or whatever we're at.

Also, next is Kathy who is now joining us as part of the SSAC secretariat type of assistant. So she will be helping us as well over the next while.

So, welcome to the ICANN 51 DNSSEC workshop. We can go to the next slide. We tried our clicker, but it didn't work this morning, so here we are.

This is the Program Committee of the people who are involved with putting this together. I'd just like to reflect for a moment that these are the folks – some of whom are here. I see Russ is here. Who else is here? I'm here. Jacques is here. Some of the other ones come and go across the time and have helped us work over this period of time and they've all assisted us. It's a good bit of work. We have a weekly call that we plan a lot of this.

I will also put the plug in, too, that if you enjoy today's workshop, keep in mind that there will be another one at ICANN 52 in Marrakesh in February.

And in just a couple of weeks, we will put out the call for proposals for that. So if you've got an idea coming out of one of the sessions that we have here and you would say, "I'd love to present this to people," or "I've got some new research, I've got a new tool, I've got a new idea, I think something would be really interesting to talk about with this group of people," please let us know. Talk to us even today and say, "Hey, I've got this idea for something that's going on." All right, let's go to the next one.

The sponsors a very critical part of what happens here, because they do one thing. They make these tickets possible. So if all of you are enjoying the fact that we are going to have lunch today and that we are able to have a time to network and socialize, this lunch is brought to you by these six companies here: Afilias, CIRA, Dyn, Microsoft, .SE, and SIDN. So these are the companies. If you see people who are here from the, thank them, because their funding this year has helped with that.

I would also say that we are in the process, too, of looking at sponsors for the next three sessions in 2015. I'm sure some of these will renew – correct, if they're in here. Yes, you will. There will undoubtedly be opportunities for others. If you would like to help show your support for the DNSSEC community and help us with these sessions, that would be greatly appreciated. Next slide, please.

One of the things that happens at these – how many people were at the implementer's gathering on Monday? A good number of people are here. One of the things we try to do is provide informal times for people to get together and network and talk. So we've met generally at a pub, restaurant, bar, somewhere in the vicinity and it's been an informal time when a group of people can sit around, have some food, have some beer and just conversation.

A whole number of great projects have come out of these events – people talking together and saying, "Oh, hey, I'm doing that." Or somebody else saying, "That's a cool idea. Why don't we do that?"

So I want to mention that this one here. We very nicely received sponsors from Comcast NBC/Universal, and the Motion Picture Association of America (MPAA) were the three companies that were brought here. Jason Livingood from Comcast played a critical role in making that all happen.

But I want to recognize these three organizations because they did help us do this. Again, we like to do these things. They've been an effective way. We'll do another one in Marrakesh. We'll be looking for sponsors there.

Also, if you're interested in attending, this one was a little bit limited in space, but we do put out news about this on the typical DNSSEC mailing list – DNSSEC-deployment, the DNSSEC [inaudible]. We put it up on my and the Deploy360 blog. We put it up on a couple of different places and kind of circulate it around. The DNSSEC community and Google+, the DNSSEC community and LinkedIn. We're in all these different places. Let's go on, please.

I also want to recognize that this session is brought to you. This is an organized activity of the SSAC, of the Security and Stability Advisory Committee. This is what allows us to have this room, have the support of Julie and now Kathy and other folks that are there. The Internet Society also provides some assistance in the form of helping with some of the publicity and having my time on some of this and other pieces around that, too. We can go on.

Our program today, as you can all see on the front side of that lunch ticket, is we've got a great number of pieces. I'm going to begin obviously with this and I've got some deployment metrics and some maps to show you.

Then from 9:00 to 10:00, we'll be in a panel that is already sitting up at the front. You can see some of the members that are there. Russ Mundy will be moderating that panel and the others who will be there talking about some of the great stuff happening with DNSSEC in North America.

Then at 10:00 I'm going to be coming up with some other folks who are here to talk about the potential impact of rolling the Root KSK, rolling the root key and what might happen there. What are some things we

need to be thinking about in terms of specifications, in terms of operations, in terms of communication, pieces around that?

That's going to be a bit of a preview of the public workshop happening tomorrow, which will be over in the Brentwood room nearby here, and that's going to be a session. On the schedule, if you look at the ICANN schedule, it says it starts at 8:00 and it goes until 18:00. We're not going to be in a room for ten hours talking about the KSK roll. We're not. It's 9:00 to 12:00. David Conrad from ICANN, the CTO, is still trying to get the webpage updated to indicate that. But if you're interested, it's 9:00 to 12:00 and it will be happening in the Brentwood room, which is right – Rick is pointing to me – right over here. That should be a good longer discussion around some of these issues. But the panel today will be a bit of a preview around that.

We'll take a short break, and then we've got a discussion around deployment and operating systems. If you look at that sheet, you've got some folks here from Microsoft, from Fedora. We've got somebody here talking a bit about OpenBSD as well. I think that's what we've got on that list for today.

Then we're going to take a lunch break. Again, lunch will be here and we'll be able to have that.

Then we're going to get into a discussion on DNS and DNSSEC monitoring. We've got a number of folks here who are going to talk about some of the services, some of the challenges they've had in implementing it. We've got at least one person who's going to talk about a new service that's being rolled out now-ish, even maybe right now today. Anyway, we'll have some good discussion around this.

And in the final part, we're going to get into a bit about DANE and e-mail services. One of the interesting aspects when we've talked about DANE and its ability to put TLSA certificates into the DNS infrastructure, one of the things that's been interesting is where it has and where it hasn't taken off.

One of the interesting aspects is e-mail, that kind of mundane thing we all use, has become one of the prime places where people have found an incredible value in this. There's some e-mail services in Germany that are marketing on their websites. They have these pages that say, "Our service is better than the others because it's secured by DNSSEC and DANE." Cool! This is awesome stuff. So we're going to talk a good bit about that – what's brought that about, how it's happening in there. Then we'll wrap it up at the very end.

That's the preview of what we've got today. You're welcome, obviously, to stay as long or as little as you want, but it's generally all a great session. This is being recorded in terms of the slides and the audio. We're not recording the video here today, but we're recording the audio. I'm recording some video and we'll see what we do with that. But it is recorded, so if you want to have people watch this later, they can go and listen to it and do that kind of thing. All right, next slide.

Let's go into some counts. Go on, please. This is where I want to have the clicker.

Some great news, on one level. If we look at just the purely TLD level of the signed top-level domains, Rick Lamb sitting back over here has this site going up there where you can see the number of signed TLDs and root. You can see this really nice chart in the number of signed TLDs.

What's that chart? What caused that jump? New gTLDs. Yes. Because all the new gTLDs need to be signed. You can tell when the program started, because – woo – all of a sudden, it rolled up there.

On one level, we have this great statistic. Over 74% - 75% - of the TLDs are signed. Now, the reality is when we dig down into the second level, there's a lot of work still to be done on that. From this level, from this perspective, we're seeing a lot of it. Next slide, please.

I want to go through some deployment maps. If you're not familiar with this, we rate things on five different stages of deployment. One is experimental, meaning that we have learned or we know that people are experimenting with DNSSEC. Announced – there's a commitment. Partial – they've signed the zone, but actually there's no DS up in the root. It's not actually being fully utilized. The DS in Root and operational is when they fully move. Next slide, please.

Here's what we've got from the database. I spent a long time on a plane trip reconciling my numbers with Rick's to make sure it all worked, but at the current time, we're showing a total of – over here in the corner – 537 signed TLDs out of a total of 731. A lot of that, the remaining 200-ish, are mostly a lot of the ccTLDs. You'll see that in a couple of pictures that are there. But the good news is we're moving a long on that. Let's go to the next slide.

Let's take a look at some of the regional deployments. Here's the big picture map showing the dark green is the fully operational TLDs. The yellowish, but it's really supposed to be a lighter green, is the DS in root. And then we go into the partial and other pieces. You'll see here a pretty good picture. Let's go into the regional.

This slide shows the bigger picture with all of the domains down here. These slides are available up on the webpage here. You can also get them from the DNSSEC deployment maps that are now hosted by the Internet Society. Previously this was run by Steve Crocker of Shinkuro. We're now hosting [inaudible]. Let's go onto the next one.

In Africa – well, the outlines don't show up too well here. But in Africa, you can see the one bright spot we had here was Tunisia up here just recently signed in September. They signed .TN and the Arabic IDN that's there. I think, Rick, that was shortly after your workshop that was there, so kudos to the ICANN training organization for getting out there and doing some of that.

The other part when you talk about where those other 200 domain TLDs that aren't signed yet, you'll see they're not there. A lot here still needs to be filled in.

One interesting spot is Morocco up there is partial. They're signed, but they don't have the DS in root yet and such, and so we're hoping that we can see if by ICANN 52 in Marrakesh, it would be awesome if we could be able to show Morocco as being operational or DS in root. Anyway, let's go on to the next one.

North America. Yeah, we're there. We're signed. It's all good. Let's go on. Next one.

In the Latin America situation, there you can see some things are signed. I put a note down here. It turned out that I had some errors in the database. Actually, I should say part of this database that generates this

map is forward-looking. It bases information on what people, what ccTLDs indicate they're going to do.

I'll be honest and say that in the transfer of the database from Shinkuro over to us at the Internet Society, I didn't notice some of the things that were set in there as more aspirations than actual measurements. So when I was reconciling data, I discovered a number of sites here, especially in Latin America, hadn't met the target dates that they are at. So these maps actually show a little bit less deployment than the ones we showed you three months ago at ICANN 50. But that's why. It's because some of the stuff – now it's absolutely set as to where things actually are. But the good news is we are seeing some growth across there. Next slide.

Asia-Pacific. Sorry, Jeff, we had Australia as operational before.

GEOFF HUSTON: That was wrong.

DAN YORK: It was wrong, I know. I know. It was a big green spot, and then I looked – no way, it's not that. Anyway, hopefully Australia will come along here somewhere. Let's go on. Next one.

Europe. The good news is since last time we met was that Croatia signed .HN and Spain signed .ES. So we had two more of the ccTLDs in Europe that were signed as well. Moving along. I think that's it from the map side.

Again, the maps are part of this project that we're doing. And we've started to take this and try to open it up a bit more. There is a GitHub repository. It doesn't have the code up yet, but it has some issues that we're tracking there and some pieces.

I would be curious. We've got this database. One of the things I'd like to do is have some kind of visualization of the generic top-level domains, because we can't really put them on a map. But if anybody out there is interested in doing some visualization work or has a student or a graduate student or somebody who would like to do some visualization work, I would be curious to talk to them because I'd like to be able to show some kind of big chart of all the TLDs – the generic ones – and be able to see that. Next slide, please.

These maps are available from the Internet Society's Deploy360 site. If you wish, there's an e-mail list to which you can subscribe and you can get these every Monday morning at about 4:00 AM Eastern time. The thing goes off and you get a little set of maps and some CSV files that you can go and play with as well. Let's go on.

I also want to mention there is a project out there called DNSSEC History Project. It is out there. It's at this URL that's there. We're trying – have been trying over several years – to compile a history and document, the history of what's happened with DNSSEC and the development that's going on out there.

This is another thing. If anybody knows of a graduate student or a student out there who's interested in the archivist type of things, we have this project that we'd love to have somebody [spend] some cycles on.

But anyway, that's out there. We love contributions and all that. Let's move on.

I think that's all I'm going to say. Then we'll just get right into the panel a little bit early. Any questions about the day before we get going? Everybody all set? Okay. We'll have coffee, tea, things.

Yes? We need a mic.

[STEPHAN LORGRAM]: My name is [Stephan Lorgram] from Microsoft. So the KSK rollover 5011 thing is tomorrow. Is there anything going on today? As far as I understood, there's going to be some private—

DAN YORK: Yes. And in fact, if you're actually looking for that, it's over in the Brentwood room right now. And Warren just went for that. You're welcome to go over to that. There is a group of vendors and others who are meeting to talk about interop today around what would be involved in some of the – from their perspective on what would be involved. If you want to head over to that, it's in the Brentwood room. And there's a group of people who are meeting in that. So that is happening today and part of it will be some of the information should come in to that report tomorrow. Yes, Jeff?

GEOFF HUSTON: Geoff Huston, APNIC. I find it somewhat amusing and weird that in your maps of where you've got cc top-level domains being signed,

comparable maps of where clients perform DNSSEC validation of names they resolve is almost the inverse.

The countries where there are a lot of DNSSEC validating resolvers don't correlate to countries which have their top-level domain signed. Isn't that weird?

DAN YORK: So the higher level of DNSSEC validation is in places that don't have their ccTLD signed.

GEOFF HUSTON: Not necessarily, exactly. And in fact, generally no.

DAN YORK: Interesting. Now, is that because of the fact that some of those countries, do they have a higher proportion of usage of, say, Google public DNS?

GEOFF HUSTON: In the African countries and in South American countries, yes. In Europe and North America and Asia, no.

DAN YORK: Interesting. Oh, we should do a little bit more—

GEOFF HUSTON: It's not just Google. It's something else.

DAN YORK: And you've got that smile on your face.

GEOFF HUSTON: No, no. I don't understand why, but—

DAN YORK: Oh, okay. You had that smile that said to me, "And I know, Dan, and I'm not going to tell you." [laughs]

GEOFF HUSTON: No. if I knew, I'd say something. But no, I don't know why.

DAN YORK: Yes, I'm sure you would. By the way, [inaudible], there are some seats along here, too. You are welcome to come up here and sit down long the table, too. You don't have to sit back there as well if you'd like to. There are spaces. We don't bite too much.

Yes, there?

CRAIG SCHWARTZ: Good morning. Craig Shwartz from the .bank registry. I'm wondering if during the course of the day they'll be some time to discuss or is it planned to be discussed what cloud service providers are doing in the way of implementing or not DNSSEC.

DAN YORK: Did this gentleman over here with the grey hair and the ponytail put you up to that or anything? Okay.

That's not specifically on the agenda. I think any number of us would love to talk to you about that, though, or to talk about that. Martin's from CloudFlare who was here at the ICANN 50 event and told us about how they were going to implement DNSSEC by the end of this year. We'll see where that gets in there. But it's coming. They've hired some good people to go and make that happen.

Any other questions? Oh, man.

MARTIN LEVY: Martin Levy, CloudFlare. Just a statistics question here. Actually, a contract question, maybe. In your stats, you mentioned that there are ccTLDs that are yet to finish the process, yet to get DS records in, yet to even sign.

However, you glossed over very quickly a phrase that said certain gTLDs were not signed. From a contractual basis, if you go back and look at the contract that the new gTLDs signed, it explicitly said two things.

IPv6 – it's early in the morning I mentioned it – and DNSSEC. Do you have the stats which then address the contractual agreement to do DNSSEC?

DAN YORK: So the new gTLDs are all signed.

MARTIN LEVY: 100%?

DAN YORK: Yes. All the ones that have come out have all been signed.

MARTIN LEVY: Okay, cool.

DAN YORK: Now, the question that I had not been able – so when I said that, there were some gTLDs. There are still some older gTLDs that have not yet signed that are still floating around out there that don't have that. I'm not going to name names, although perhaps I should, but some of the older ones of the original 22 are still not signed and non-contractual. [They don't have that.] But all the new gTLDs do that.

Now, one interesting thing that I had not yet figured out a way to do programmatically is, with the new gTLDs, they have to be signed – so they have to put a DS in root. There's some vagueness in the Registrar Accreditation Agreement, the 2013 RAA, as to whether they have to be accepting – there's some vagueness in there.

But they do have to have a DS in root. They have to be signed. Whether they go to that next what we call operational where they're accepting DS records – or DNS keys – from the registrants, that's not entirely clear. But they are signed out there and they are doing that.

The other interesting aspect is they all have to contribute their records to the centralized zone database (the CZDS) so you wind up being able to have some interesting stats out there.

One of the sites that I found useful is NTLD (new TLD stats) ntldstats.com, and they will show you right up there how many are signed for each zone of the new gTLDs. It's currently about 1.44% of all of the three million new gTLD domains.

MARTIN LEVY: This sounds very much in line, the second order issue putting an old hat on, on the v6 side is that, although the language talked about v6 for name servers and everything else, turns out WHOIS servers type kind of got forgotten. And then the subtlety. So this is all second order type issue.

DAN YORK: Right, they're subtleties.

MARTIN LEVY: Thank you for the clarification.

DAN YORK: Yeah, there are some subtleties in there, but a good number of the sites are doing that. The one interesting question I just haven't yet been able to figure out for these maps is exactly how to find out of the 400 new gTLDs how many of them have started [inaudible] signed registrations

from registrars and others that are out there. So, need to figure that out some way. Anyway, it is happening. Any other questions?

All right. Well, I will yield my five minutes to the next panel, and we'll start out with our session on DNSSEC in North America.

RUSS MUNDY: I'm Russ Mundy, and let me add my welcome to Dan's to everybody in the room here. We're thrilled that you could join us. As we normally do for these workshops, we try to feature things going on in the geographic area where we happen to be holding the meeting. Since this one is in North America, it's labeled North America. That's who we have as a group of panelists here to help us talk about what is going on here in North America.

I think we are following the order on the agenda. Paul Ebersman is the first of our panelists and he's going to give us some insight as to what Comcast has been doing with the DNSSEC realm.

PAUL EBERSMAN: Thank you. Go ahead and advance. I'll talk a little bit about the history. I can't take credit for that. That was before I came in. Next slide.

The wonderful summer of 2008, which anybody who was involved with DNS probably remembers all too vividly. That was sort of the first sign that rolling out DNSSEC was going to be useful since cache poisoning was out in the wild.

At that point Comcast started beating on all of their vendors, because in 2008 DNSSEC was still very early in vendor support, key rollover, any of the tools at all.

It actually took several years to go through. They got most of this stuff done, ready on the infrastructure by 2010, but since at that point also there weren't various useful things like negative trust anchors, people were unsure of exactly what was going on, they weren't going to sure [inaudible] with customers, they did a fairly conservative rollout over about a year or 15 month until it was pretty much fully deployed for all of the recursive customers. I'm not sure exactly what it was then. Current count I believe were about 22 million customers.

They also started going through [and] signing zones, and by 2012, all of the Comcast-owned domains were there. We have since progressed to being able to DNSSEC sign if we are running the authoritative zone for a customer. We are in the midst of rolling up NBC/Universal which was bought a couple of years ago and getting them DNSSEC signed. So at this point, everything externally visible is signed and validated. Next slide, please.

There were a number of useful and painful lessons that were learned in the early deployment. There was a lot of hardware and software that had to be upgraded to deal with it, particularly some of the signing. There were some very old machines lying around loose that when we had static zones didn't really need to worry about it, but suddenly when we were ramming lots of things in and signing and resigning. Got quite painful.

We discovered all of the things that people we thought knew about DNS but the fact that [TCP] is not just for zone transfers that what happens if you do have a packet over 512, did they handle [inaudible], do they handle packet fragmentation? So, huge education for our network and security folks. And then obviously beefing up the authoritative infrastructure. It was painful in and of itself.

The other interesting exercise was that under cable.comcast.com and comcast.com, there were a number of groups that had been delegated sub-zones or sub-domains. So the mandate was if you were under one of the signed zones, you will sign as well.

So they first had to go out and find out who actually owned all of those name servers, since that wasn't as well recorded as perhaps it should've been, and then getting all of those folks to understand that, yes, this was mandatory, what the hell was this, how to do this and all of the support. That also took I think pretty much most of the year. We still occasionally have that where we have some new group that [has decided] that they want to control their own destiny and their own zones and they discover that some shiny DNS server they thought was really cool does not do DNSSEC. So it's sort of ongoing. Next slide, please.

One of the things that we learned, because we love playing with multiple moving parts at once, we were in the midst of also rolling out a v6 upgrade and it turns out that, even though it is more operationally complex in terms of planning, if you're going to upgrade anyway, upgrading and mandating that you must support both v6 and DNSSEC

and beating on your vendors for both at the same time wound up ultimately being a cost savings.

We did also have an awful lot of work to get our first-year folks to understand again beyond even that it was DNS what DNSSEC was, why certain things worked for some people and failed for others, monitoring which – I'm going to have some questions also when we get to that session. It's still an ongoing – what are the key metrics that you look at to figure out if DNSSEC is actually working reasonably well?

We are still suffering with the registrar and automated, or at least less clunky methods of updating your DS records when you do KSK rollovers. Next slide, please.

So now we'll get into my tenure of having come into Comcast. Next slide. We are fairly comfortable at this point. The first-year folks have learned to at least do basic things, and I'll go over in a slide upcoming what the process is, but they're actually looking at failures and they are thinking of the fact of, "Oh, let me try doing it with checking disabled to see if that works," so they can escalate.

Most of the things that we are seeing are the usual suspects. Most people when they initially sign their zone do it correctly. The problem is ongoing. Do they let their zone signing keys expire or signatures expire? Do they update their KSK, but forget to do the DS? Does someone mistakenly delete the DS and leave the zone signed? Those are all the things.

Not to be too cruel on naming names, but at least in the U.S., many of ours seem to be folks in .gov and .mil. We do have a process for putting

in a negative trust anchor where, at a certain point in the validation chain, even though it fails validation, you essentially ignore that and give the customer the data anyway to allow people who have broken their DNSSEC to still be reachable by our customers.

One of the things we've also found both for our support folks and for the folks who have their zones broken to explain it, DNSViz is a wonderful tool. I'm thrilled that [Casey] has started doing a release where you can run it internally, but right now it is available on the web. It is a very good graphical breakdown of exactly what goes through and it's been very helpful explaining to people who did DNSSEC but didn't quite get it what went wrong. Next slide, please.

So we are fairly careful. We actually have our first-tier notice or we have some monitoring that can check it for our own zones. They're supposed to actually verify it themselves, run DNSVis, and then they escalate to us. It actually escalates to our senior engineering staff, and one of the things that I like about Comcast is the fact that they have stepped up to saying it is acceptable that one of the costs of doing DNSSEC validation is the cost of having mid-level to senior DNS engineers actually calling domain owners on the other end and working with them in order to get things fixed.

That doesn't always work. Particularly with .mil it's very hard to get a contact who actually owns that subdomain. [Scott Rose] is a wonderful resource, but sometimes he goes on vacation or goes to sleep.

So when we have something like that or we have a very critical one, like when whitehouse.gov suddenly falls off the map, we will, at that point [inaudible]. So far that has been a lot less prevalent than I thought.

We're doing one about every 6-8 weeks. In most cases they don't stay on for more than a week or two. I don't think we've ever had more than two dozen in total. I think we're about 21 NTAs right now for the entire Internet, of which six are essentially one division who is broken into multiple domains.

It's not a panacea, but it's certainly valuable and it does allow us to actually keep DNSSEC running without having to disable it for everything. And that's it.

RUSS MUNDY: Paul, thank you very much. As always, very interesting presentation. Comcast has done a great job with getting DNSSEC out there. I know personally I use you guys as a great example of how any large, competent, capable ISP can put this out and make it work. You guys have done a superb job with that. And I'd like to hold questions until the end of the panel if we could.

And next we'll go on to Jacques Latour from CIRA who is going to give us some insight about what they've been doing for .CA.

JACQUES LATOUR: Hi. Well, I guess everybody knows I'm Jaques. My name is Jacques Latour. I'm the CTO with CIRA, and we're responsible for .CA domain in Canada. Next slide.

So we started to work on DNSSEC in 2000 and something – a long time – and we had multiple iterations of our deployment of DNSSEC. And

finally in 2011-2012, we decided to put a lot of effort behind it. It was basically a three-phase project.

Phase one was to sign the zone and we did that in January 2013. We have something that is somewhat unique – Dual in-line signer. I'll cover that quickly. And it works great.

Phase two was to implement DNSSEC in the registry and we did it the Canadian way, a little bit different than everybody else. But I'll cover that, too.

And right now we're focusing on promoting adoption and I think that's the area that needs a lot of work for DNSSEC at this point. Next slide.

So basically our DNSSEC signer solution is a bit different than everybody else in the fact that we sign the zone twice. Once with different technology. And we sign with BIND, for example, and we sign with OpenDNSSEC and we compare the output of both zone file and we make sure there's no bugs that were introduced in either one. If they don't match, we stop the whole thing.

And we have a production and back-up site. Both are running in parallel and in production. So it's working great. We don't have a lot of issues. It's been one year and nine months in production. We did 21 ZSK rollover, and last January we did KSK rollover. We have a DPS, or DNSSEC Practice Statement, that states how we do key rollover and all that stuff. We just did it and it worked great. No issues. January next year we'll do the same thing. Next slide.

The second phase was to put .CA in the DNSSEC in the registry. And the objective was to keep it simple for registrars. I think in the last year we

talked a lot about some registries supporting DS, some supporting DNSKey. Next slide.

We decided to do both. We support the entire RFC. So there was a provision in the RFC to support DS or the DNSKey or the DS and the DNSKey and we support all of that. So we can work with any registrar. They do whatever they want and we don't do any politics around if we're DS or DNSSEC. And I guess it works. Next slide.

The adoption. There's three circles and there's different ecosystems that we need to work with. Right now we're focusing on the getting the ISPs to validate. We also need to get hosting provider to do DNSSEC. So those are different areas that we're working on. But right now our focus is on getting the registrars to support DNSSEC and also working with the ISPs in Canada to validate. So that's our strategy [inaudible]. Next slide.

So if you go to CIRA.ca website, marketplace, you can see which registrars support DNSSEC for CIRA, for .CA. They all support the Web interface. So with our registry, [inaudible], we have a Web interface where you can do all the EPP stuff in there basically.

All of our registrars today support DNSSEC Web. They don't do EPP. So when I said [I don't know if our] EPP and stuff works because nobody's using it. It's been tested, [inaudible] and all that, according to spec. But I'd love to have a registrar actually do a DNSSEC provision. So that's our goal: to get them to do it.

And we have 93 signed domains. Yeah. I don't know what to say about that. [laughter] It's more than one – there you go. Next slide.

So thanks to Jeff, we have good statistics now in Canada on IPv6 adoption, on DNSSEC adoption. It's great stuff. So I'm going to be using a lot of this in the upcoming weeks. We have a bunch of conferences in Canada, mostly related around ISPs and I'll be pushing this stuff in their face.

Instead of being negative on focusing on the ISPs that don't do their work, I'm going to outline or tell and get people to stand up on the ISPs that actually do validation and actually do IPv6 and all that. This is a great tool for this. I'll try to not [change] them, but create some competition I guess. Next ICANN I'll report on how that went I guess. Next slide.

So, basically if you have any questions, contact us. If you want to talk about KSK rollover, Jake, Zack there, raise your hand. [inaudible] planned that and it worked flawlessly. [inaudible] manager. So if you want to do EPP, he knows how to do that stuff. And that's it. thank you.

RUSS MUNDY: Thanks very much, Jacques. I appreciate that. Now we'll go on. Scott Rose is our next presenter from NIST and Scott's going to give us some insight as to what's going on in the .gov realm.

SCOTT ROSE: Hello, everybody. I'm Scott Rose from NIST. Before I get started, I'll be using one acronym a lot, probably. It's FISMA. It stands for the Federal Information Security Management Act. It's an attempt to kind of have a government-wide [inaudible] policy and framework. Often that's what agencies are audited against, as if they were individual companies. But

all agencies are independent on how they meet these requirements, but the requirements are supposed to be standardized. You can go on to the next slide.

This is what it looks like today. Actually, that pie chart is probably about a couple weeks old, but I'm guessing that number hasn't changed. We're holding steady about 82% signed. This is federal executive branch. A lot of the ones that aren't signed are legislative and judicial. The courts and the Congress, they don't have to do FISMA, so some of them aren't because it's just the way that the federal government works.

Also, the federal government, .gov, also includes states and locals. So the bottom bullet there, .gov is about I think 5,000 or 6,000. I'm not 100% sure. Of those, at the beginning in 2008, about 1,700 of which were federal. That number has now shrunk down to 1,393. That's a second order of thing that started to happen is when you're told to start doing DNSSEC, a lot of agencies looked at all the delegations they had and said let's start consolidating. Let's get rid of some of these. We don't have to do them. So that's what we've seen. We've seen a lot of shrinking of the number of delegations.

This number started out really low in 2008. That's when the mandate to sign .gov went out and is also included in FISMA. It started to – it was growing sort of steadily until about 2011. That's when what called the DNSSEC Tiger Team was formed. It was actually co-chaired by Department of Homeland Security and NIST. The goal was to figure out why deployment was so slow and speed it up.

Once that started meeting and started to set up a compliance-checking program as well as just getting a list of all the admins, that number shot up to about 80% pretty quickly, as well as the number of errors. In the first few years, there would be 100 signed delegations or a couple hundred, 10% of which at any given time were bogus. They had some sort of error.

The length of time it took for them to fix that stretched on for weeks or months. After the compliance checking started, we saw a lot fewer errors, down to 1 or 2% and the time it took to fix them until they were resolved shrunk to days – sometimes even one day. Next slide.

So lessons we learned over the years. The first one is set up for a community of interest. If you have a TLD, if you have a large organization, if you're just a bunch of companies that decided you all wanted to do this, set up a monitoring regime, a compliance checker to make sure everybody's still doing it. This is handled right now by the Department of Homeland Security, their operational side of the house which sends out queries to every agency looking for DNSSEC, looking for keys, checking validation. Then it sends out a weekly report as to how everybody's doing.

These go to the federal CIOs as well, so the CIO gets a report saying, "Your agency is not in compliance," or "Your agency is in compliance." And they get taken to task for that when they meet with the whole e-government [head]. It's just like IPv6.

The biggest one is also ensure every organization has an updated list of points of contact. In the beginning, you often didn't know who was in charge of the zone. And there's a difference between the administrative

contact, which was usually a federal employee who had no technical experience, and the technical contact who was a contractor who actually did the day-to-day operations. And if that contract ever changed, sometimes that point of contact wouldn't be updated and then things got wrong really fast.

We also encouraged automation. And this was another second order effect. We saw two things that started happening in agencies. One, to move to automation. There used to be a lot of running BIND in-house. A lot of that moved now to [appliances]. They said, "Well, I'm not going to bother training my staff or the contractors who are going to change every one or two years. I'm just going to buy product X and it's going to handle it all for me." So we saw a huge uptake in these [appliances] that did signing.

The second one was these consolidations. Not just getting rid of delegations, but in some agencies – for instance NIST and NOAA – they're the weather people (noaa.gov) – they actually took all the delegations that were operated by individual units or divisions that are spread across the country, they allowed them to maintain their zones, but the actual signing and key management was done by one group in DC in their IT department. So they would actually act as kind of like a signer bump in the middle.

They would transfer all the zones up, they'd sign it and then they'd push it back out – send it for actual hosting to the authoritative servers which were maintained by these individual units. But they did all the signing and all the key rollover. That way they could essentially manage everything and try to reduce the number of errors.

One thing that we tried to do – we weren't 100% successful – was to kind of foster this community of admins. We had an e-mail group. Because there's a lot of gov-centric questions that a lot of administrators didn't feel comfortable with asking external communities, external [lists], because they were told not they're not to discuss their internal operations, as well as you're kind of frowned upon to talk about actual products, if you have an e-mail address ending in .gov, because you're not supposed to be endorsing anything.

So we had a closed membership list, just for either government employees and contractors directly supporting a government deployment. That was really helpful. Next slide. I think it's the last one.

Some new issues coming up. Now FISMA has been updated. Every couple years, it's updated. Now we're requiring DNSSEC validation. That is actually kicked in right now. So agencies are supposed to be deploying validation. They're supposed to turn it on. Originally they didn't have to. They just had to sign. Now they have to sign and validate.

Some agencies already are. Big one being the Department of Treasury, I believe, which is kind of important because they do money. Homeland Security is still doing compliance checks. There's been a rollover in personnel, so a lot of this is on dead reckoning. It's just, "I know this tool. I'm going to run this tool. I'm going to generation this report." They're not doing a lot more helping like they used to because the new staff isn't quite up to speed on DNSSEC.

A lot of the errors we're still seeing are due to, I believe, administration turnover. Contractors leave. Full-time employees leave, and they don't write down their procedures. So when it comes to do something, it's a

bunch of new people and they're making the same mistakes the last group did because they didn't have a chance to learn. We're especially seeing that with KSK rollovers. That's still the biggest issue. Even in states, too, we're seeing some deployment in the states.

Last one, agencies are being also urged to what they call move to the cloud. So when they do, they ask their cloud provider, "Well, I have to meet all these FISMA requirements." And oftentimes the provider doesn't do DNSSEC. And then you get into an interesting confrontation where "I have to move to the cloud" because that's one mandate we're supposed to be urging. And we're also supposed to do DNSSEC and FISMA. And that's another mandate we have to meet. So which one is more important?

I said sometimes you get the ugly thing where agencies want to move to the cloud, so they use DNSSEC looking at a DNSSEC problem as that's their excuse to not do it. It's not because the cloud provider doesn't do it. It's because, "Well, I've been told DNSSEC is bad, so we're not going to do it. We're just going to move to this cloud and handle it that way." That's usually the wrong answer, but that's what's given.

So we're seeing some of that, and a lot of it is just "Well, we're using this anyway, so we know we're not out of compliance, but we don't care." That's the other one. And if you're a big enough agency, you can get around that. You can get away with that answer. Smaller ones usually don't.

That's about it. Like I say, we're seeing increased uptake in the states, which is nice. So we'll see how well validation works in this coming year, because it just kicked in, so the actual auditing won't take place until

probably this year. Then hopefully we'll get some numbers as to the validation deployment levels.

RUSS MUNDY: Okay. Thank you very much, Scott. We'll do the questions at the end, so if you don't mind waiting. We should have plenty of time, though. Next we have Duane Wessels from Verisign. Verisign has been a very strong supporter of DNSSEC activities here. Duane is one of their key people for the lots of good things that we're seeing come out of Verisign. Duane?

DUANE WESSELS: All right. Thanks, Russ. So I have a few highly-graphical slides to talk about some trends that we've been able to discover in DNSSEC registration among COM/NET registrars. Next slide, please. By the way, we don't have a laser pointer do we, by chance? Okay, that's fine.

This graph is probably familiar to a lot of you. This shows the number of counts of signed domains in .com and .net going back to late 2010. Just FYI here, what we're calling – sometimes what I call a signed domain is really just a domain with a DS record. I'm not actually validating that it's signed properly. This is interesting, but we can do better than this. Let's dive in a little bit deeper. Next slide.

So one of the reasons that we are interested in this is we want to know if DNSSEC option is being driven by the registrars, perhaps, or is it more driven by the users or maybe none of the above. We're interested in knowing are some registrars, do they have a significantly-higher

percentage of signed registrations and how have those numbers been changing over time? Next please.

So this chart shows, again, for .com and .net, it shows on the Y-axis is logarithmically scaled the number of signed registrations in those zones. And on the X-axis, it just ranks them by that value. So you can see that, for example, the registrar that has the most has somewhere above 100,000 signed registrations and it's your typical Internet-like heavy-tailed distribution tailing off. I only show the first 50 here, but I don't actually remember how many total registrars there are with signed domains. It's in the order of a few hundred. And the pattern for .net is more or less similar to for .com. Next please.

So this scatterplot shows the percent of signed registrations for each registrar. Each dot represents one registrar and the X-axis – again, logarithmically scaled – shows their total number of registrations and the Y-axis shows the percent of those that are signed with DNSSEC. So there's a couple that sort of stand out at the top, but most of them are way down on the bottom there and we can't really differentiate them.

So on the next graph, we show the same data, which is just scaled logarithmically again. The first thing you might notice is this short of sharp diagonal line. That just represents those registrars that have just one signed registration. So the one that are more interesting, of course, are a little bit higher up and to the right there. It's sort of a nice distribution, but again with the logarithmic scaling of the Y-axis, you can see that most of them are down in the range of about .1% of signed registrations. Next, please.

So looking over time, this graph shows kind of the same thing, but going back a couple years. Here each line or each color is one particular registrar, and this shows how the counts have changed over time so you can see that some of them have pretty aggressive adoption of DNSSEC. Very strong up and to the right. Others are less so.

This blue one is kind of interesting. So they were very [inaudible] for a while and then it topped off, either because they ran out of registrations to sign or something changed and they stopped signing their new customers. I don't know.

The last slide is a little bit more – the last one shows the same thing, but again by percent of signed domains. One that's very important here is this purple one at the top, which starts out with something like 90% of all their registrations signed and then starts to taper off. This would seem to me to be some kind of change in policy or procedure or something where they were getting new customers but not signing them up by default or something like that.

There's this blue one which we talked about before. You can see the very aggressive growth. And then it tops off – it flattens out at something like 85%.

This black one, which starts off at 100% and drops down, I didn't really investigate that. The 100% could be a very small number. Maybe even just one. I don't know how many total registrations they have.

Then down at the bottom you can see this green and this red that are sort of interesting. One of those is – they're both almost flat, but one is slightly less flat than the other. Those are interesting also.

Lastly, I just want to make a little pitch for Verisign's cloud-signing service which is offered free to registrars. We certainly hope that more and more will take advantage of this. If you are listening or reading and would like more information, you can send e-mail to that address listed there. If you're a geek like a lot of us in the room, you can even go and read the DNSSEC Practice Statement and maybe learn a little bit more about how the cloud-signing service works. So, thank you.

RUSS MUNDY: Thanks, Duane. We have I think got a very good broad coverage here. I saw some questions, hands being raised, as we went along. But we wanted to hold the questions to the end and now is everybody's opportunity to raise questions. So, please, for any of the individual panelists or a general question for a panel. Who wants to be first? Let's take that gentleman over here.

UNIDENTIFIED MALE: Hello, this is [inaudible] from Microsoft. I wanted to know that last slide [from Verisign], can we make any [inaudible] of data from that? There are things that are increasing, sudden falls. So can we just predict something? Is there a particular type of registrars which are going to sign or are there a particular type of people who are not going to sign, phasing out DNSSEC or something? Can we predict a trend or [inaudible] patterns here?

DUANE WESSEL: I guess you kind of can, but I'm a little bit skeptical because, on this graph, the ones that are sort of interesting [inaudible] here. There's

maybe five or six registrars out of hundreds that are taking registrations and maybe even signing.

I think just from this data it would be wrong to draw conclusions about the whole – long-term conclusions, because it's such a small sample.

UNIDENTIFIED MALE: Or [inaudible] that may reframe this. Is this a behavior or is it a type of registrar which behaves in a particular way when it comes to DNSSEC? Can we establish that at least?

DUANE WESSELS: I think you can, yeah. But again, we're only seeing that behavior for a small number.

RUSS MUNDY: Dan, go ahead.

DAN YORK: So thank you to all the panelists for some great presentations and insights there. Jacques, with your 93 signed domains, I was caught by the fact that last night at the music night there was this air hockey table that was occupied by the Czechs and the Canadians who were having a dueling battle of who could win the most things that were there. I have to say, though, that the Czechs are clearly winning in the signed domains, the .CZ domain. So I don't know. But anyway.

I wanted to go to Paul Ebersman. You guys at Comcast, thank you. You've been doing great work over the years. One question I had: in the

past, you were doing some work with notifying the public around some of the domain failures you were seeing and you had, at one point, a Twitter account. I don't know if that's still active. You had some ways that you were doing. Is that something you're thinking of still doing, or are you still doing it? I haven't really paid attention, to be honest.

PAUL EBERSMAN: If it's the page that I'm thinking of, we've been in the midst of upgrading a bunch of servers and webpages. We've had a couple of re-orgs and other things have gone through. So that was stuff that was running before I even got there and ceased before I got there. I think it hasn't been running since last November. So assuming our lawyers don't have any issues with it, I have no problem.

We have, for high profile, announced them via more common Comcast Twitter accounts. We occasionally get the conspiracy theory. There was one where there was a gaming site that failed validation and we were getting this whole fling-fest that "Comcast is preventing me from playing online poker."

DAN YORK: Cool!

PAUL EBERSMAN: Yes. So we basically had to do the "No, no, no. They chose to do DNSSEC signing. They blew their key rollover. You can't get there from us or Google and that's why." But yeah, it's mostly been on an ad hoc basis

for more high visibility. At some point, I would like to get an automated "here's who we see failing."

DAN YORK: Cool. Okay. Well, thank you. I know it was something that [inaudible] before you, so it's not really fair to ask you some of that.

Scott, I had a question for you, and my last question is when you talked about cloud providers and the challenge they have moving to that, when you say cloud providers, is it DNS hosting providers? Is it CDNs? What are the kind of providers that they're moving to that are the barriers because they don't support DNSSEC?

SCOTT ROSE: It's mostly the CDNs. I should probably clarify. The hosting side seems to be well underway. Or if they're not doing it already, they usually can do it quite easily. It's like the hosting side of a cloud provider. They'll usually be able to do DNSSEC. In fact, we got a lot of agencies that are customers of large providers and they're all signed. It's usually on the CDN side, because they've got webpages that they need to have up. That's where there's usually less DNSSEC. So it's usually that C name out of the .gov to the CDN, that's the last signed response you get; and then from then on, there's no DNSSEC.

There's a lot of internal policy debate about whether or not that counts as meeting the compliance or not, because the agency response is, "Well, everything in the .gov is signed. We have no control over .com and .net." And then there's been some, obviously, the security professionals pushback saying, "Well, that's not the point."

If an agency is going to a provider for hosting, they're usually okay, depending on the hosting service, but the majority of them want to be able to say they can meet FISMA compliance because they want to get the customers. It's usually the CDN stuff where they run into trouble.

DAN YORK: Thank you.

CHRIS BOYER: Hi, I'm Chris Boyer from AT&T, but I'm actually asking this question in a different capacity. I serve on the NIST ISPAB, so I was going to ask you a question about whether or not DNSSEC is built into the FedRAMP requirements for onboarding cloud providers.

SCOTT ROSE: In earlier versions, it was. It actually was a line item that you had to do DNSSEC. I don't know about current versions, but usually if you take the broadest net of FISMA, if you are a system that is possessing, hosting, or processing federal information, you would fall under it.

Depending on the current mood at the Office of Management and Budget, they like to bring contractors under that umbrella, and other times they're not. Sometimes they even want to bring universities under that. If you're a research department at a university and you're getting data from the CDC or NIH, there's some questions whether or not you fall under FISMA. It's depending on their current mood.

CHRIS BOYER: Yeah. I was just asking because for folks, if you're not familiar with FedRAMP, it's basically how cloud providers are being onboarded to be offering services to government agencies. I was just curious if DNSSEC is an issue for the cloud community. If it's not something that was built into the FedRAMP requirements, I don't know that they would be paid as much attention to it as otherwise, right?

SCOTT ROSE: Yeah. I remember in the first versions of requirements that went out, it was actually explicitly called out, which was rather strange for something that's usually that low level. Later versions, I don't know if it was just kind of buried under the other security stuff, that kind of stuff. So yeah, I'm not exactly sure of the current state.

DAN YORK: Yeah. Please, go ahead.

UNIDENTIFIED MALE: [inaudible] from New Zealand Registry Services. I have a couple of questions, one for Jacques. Do you see any concentration for a single registrar in terms of DNS support or signed delegations?

JACQUES LATOUR: No. Most of it is ad hoc across our registrars. They do it manually through our web interface. We don't have a single large registrar or hosting company that's signed all their domains yet. That's what we're trying to work on to increase numbers.

UNIDENTIFIED MALE: Okay. Because we have – [inaudible] we have DNSSEC support and signed delegation 150, and I will say like 110 of them come from a single registrar.

JACQUES LATOUR: [inaudible] to run a query.

UNIDENTIFIED MALE: Yeah, cool. And do you have any registrars doing signing, DNSSEC signing? No?

JACQUES LATOUR: No. That's public. We deal with GoDaddy, for example, and they haven't built the – they haven't done DNSSEC accreditation. So that's [the thing]. We have a process [inaudible] with DNSSEC and people are doing it, but we need more awareness and education in Canada on how to do it.

I have a question for Duane. Can you share who your largest signed registrar is or that's private?

DUANE WESSELS: We decided not to name names, at least not at this time. So, no, sorry.

UNIDENTIFIED MALE: I have one more short question for Duane. Your free signing services, do you have any users at the moment?

DUANE WESSELS: Yes.

UNIDENTIFIED MALE: Do you have any idea how many domains are being signed using that?

DUANE WESSELS: That I don't know. I can try to find out for you.

UNIDENTIFIED MALE: Okay, cool. Thank you.

UNIDENTIFIED MALE: Jacques, could I just ask about your accreditation. You said you have an accreditation process for registrars to submit DNSSEC records. Is that it?

JACQUES LATOUR: [Stuart] can talk about it.

[STUART]: So we actually have two levels of accreditation. So the first is before someone comes on board just for the Web piece. It's a fairly straightforward, more of a questionnaire to make sure that they have a level of understanding. It's more of an education piece than an accreditation to try and do a bit of knowledge transfer with the registrar and their support division to explain the difference between DS, DNS Key, how you use them, where you get them, who you get them from, all those kinds of things if someone wants a signed domain.

And then for EPP, we have an accreditation process for EPP anyways, and so we've added an extra suite of tests just around signing a domain through EPP and with our extensions, which is basically the base extension but just in our – because we accept DS and DNS Key and adding and removing keys and all that kind of stuff.

But, unfortunately, no one has actually done the EPP accreditation yet. We have a couple on queue for that.

DAN YORK: We have a question back here.

[MICHAEL]: [Michael] [inaudible] Amazon Web Services. I have a question for Scott and the rest of the panel. How much do you guys care about NSEC3 support?

SCOTT ROSE: How much do they care about NSEC3 support? Depending on who you ask, it's either very, very important or not important. Some do want the [inaudible] a little bit, [so] they do deploy NSEC3. I noticed that a lot of the appliances that come out that agencies are buying turn on NSEC3 by default.

So I would say probably the majority of agencies that are deploying are doing NSEC3. There's a few that are still doing NSEC and they've just – because they have a public statement of saying, "We don't care of you lock our zone because it's publicly facing anyway. There's nothing stopping you."

But those that are kind of aware of DNSSEC, they tend to want NSEC3. I know .gov itself, the actual TLD, uses NSEC3. So I'd say, eh, those that are aware tend to favor it and those who are not – and those that are just buying something off the shelf just go with the defaults.

UNIDENTIFIED MALE: Thank you.

DAN YORK: Julie, we have a question in the Jabber room.

[JULIE]: Now it's gone.

DAN YORK: It's gone, okay.

JACOB ZACK: Jacob Zack from .CA. My question is for Paul from Comcast. Again, I want to thank Comcast for taking the lead on DNSSEC and being very transparent and open in everything they've done so far.

You mentioned the negative trust anchors and how they are put into the zone. My question, two parts, I guess. To remove a trust anchor, maybe you can tell us, is it an automated process where you're constantly monitoring that broken domain to instantly put it back? Is it a time basting where you check every six hours?

PAUL EBERSMAN: We have some [inaudible] checks that will kick off in [inaudible]. I don't remember the interval. Once that happens, though, we then do a manual process anyway. We're actually in the midst of resurrecting the negative trust anchor draft in the IETF with some recommendations.

One of the interesting failure modes is if one of your authoritative servers is still [borked] while all of the others work. So you actually need to go through and make sure that everything is serving the same data and that it is all validating. That's [actually] more complicated and we haven't had the programming cycles to do a better check. But currently we will remove it only manually when we make sure that it won't just fall over and fail again.

JACOB ZACK: And the second part to the question. Because you have such a large base of recursive information to be able to determine these broken domains, would it be possible, perhaps only helping those outside of the U.S. – the competitive base – to provide a real-time feed for, say, a cable provider in the Netherlands. They may be facing the same issues, or maybe they don't want to implement DNSSEC because they're worried about "it will be broken for us but not for our competitor."

If there were a real-time feed mechanism out there, would it help to expand DNSSEC adoption?

PAUL EBERSMAN: I'm not sure if it would help to expand it necessarily. I do know that, because of the fact that we are also providing phone service and a bunch of other services, we have a bunch of FCC requirements. So

sharing data for us is one of those [trip to] the lawyers and the data review boards.

We could probably scrape it down to something that would be easier to get out publicly, but right now, for better or for worse, we're actually throwing engineers at the problem, and when we notice failures, we're actually trying to reach out to them and educate them in the hopes that they won't get an automated response every time they screw up their zone signing key again. Hopefully we educate them and they don't have that problem again.

DAN YORK: [inaudible] more questions. We've got such a great group panel. Oh, Jeff. Thank you.

[JEFF DUNN]: Since we've got both ISPs and registrars in the room, is the onus and responsibility, if there is one – I'm not even sure there is – but if there is such a role of [inaudible] validating what you see in DNSSEC and alerting those folk that there's a problem. It's not really the resolver side of the world or actually the DNS registration side of the world. I'm interested in views and opinions, because after hearing someone say, "I want a feed from the ISP side of…" basically domains that are dead that, for some reason, shouldn't be, what about the registrars? Don't they do this? Shouldn't they do this? What's their role?

UNIDENTIFIED MALE: Well, I'll take first stab in that part of the problem with the ISP feed [as] the method is until one of my customers actually cares and tries to reach a dead zone, we don't know about it, so that's already an incomplete data set.

That said, my personal opinion on the registrar side is having that potentially as a value-added service that you opt into to be monitored would certainly be something I'd like to see, mandating that the registrars must watch everyone I think is something that will pretty much do [inaudible] because I'm not sure they'd buy off on it.

UNIDENTIFIED MALE: Yeah. I'm not a registrar or a registry or an ISP operator, but I do have an opinion. Coming from a community where reputation is a lot, we've been trying to kind of get that into the registrar world, because like I said, there could be some small national park that no one ever seems to go to but has a DNSSEC failure and one day it gets noticed.

But for our side, it's much easier. We have one registrar and we have one group that's tasked for checking everything [inaudible] where it should be done. We're trying to get more proactive in that, trying to get the agencies that are responsible for that saying, "Let's not just check, and at the end of the week, give you a notice that you failed," but kind of say, "Look, we've noticed your signatures are going to expire tomorrow," or your key hasn't been rolled or we see a new key, that sort of thing – trying to get a little bit more proactive in alerting agencies that something may happen rather than something that has happened.

[PAUL EBERSMAN]: So I have an opinion, too, but I want to be clear this is not the company's opinion. This is more from personal experience that one of the frustrating things about – I've tried to build a DNSSEC [inaudible] thing at least once, and the problem you run into is, with signatures, you don't know when they're going to roll it. You only know that they get closer and closer to not – something's expiring.

So it would be good if the community could maybe develop standards that says, "For the benefit of everyone, always do your rolls or whatever at least X hours or X days before they go bad," so that you have some lead time. Otherwise, if you develop an alerting system, you're always going to get false positives.

UNIDENTIFIED MALE: Actually, if you haven't been giving feedback into the draft, I think [inaudible] is the one who's doing it right now. We're actually trying to come up with an operational set of recommendations for DNS standards and rollovers.

DAN YORK: And I'd like to add just a little bit more to the responses to that question. Great question, Jeff. It is a real challenge to sometimes determine where the responsibility is most appropriately placed for things of this nature.

There's a panel later on today dealing with DNSSEC monitoring, so you'll see some capabilities from a monitoring perspective, but that doesn't

actually fix things. It can give people hand waves and say, “Look out! Something’s coming!” But it doesn’t actually get things fixed.

And one of the problems that I think we have – as a result of the wonderful great flexibility that DNS has always had as far as how it can explicitly be done, you use the term if it should be the registrars or the ISPs.

Well, from the perspective of the ISPs, it’s a business issue if their customers can’t get to places where they want to go. They’re concerned, but it’s not really their responsibility to fix it, but they have strong motivation to.

From the other side of things, it may not be the registrar that’s operating the name service. It’s really whoever is operating that name server and doing the signing that should be responsible, and we do definitely need a better way to figure out how to even describe it properly so people understand what we’re talking about, because most of the names in use in the Internet have a name server that’s provided by the registrar. It’s far from 100%.

So getting that allocation of responsibility in place and having qualified people to do it is important.

[JEFF DUNN]: If I could just add one more. The problem from the ISPs is, and sort of the outside observers, we can’t enumerate zones [inaudible], so there’s no way that anything we do is anything other than the domain names people use.

Automating that process of going, "Well, they're using it a lot and I think I'll white list it," runs into an engineering problem – a social engineering problem – that you might try white listing the [lie].

And at the moment, it's sort of inconsequential, but as soon as you start talking about DANE, white listing the [lie] has disastrous implications.

And I must admit, from the perspective of my side of the fence, I am not a registrar or anything like it, I would like to see the registrars be a little bit more proactive in detecting these kinds of things, because in theory, they know the difference between what is true and a mistake. I'm serving good data. If you're seeing bad data, this is what DNSSEC was meant to do.

DAN YORK: Well, and in fact, early on in the specification developments, there were several times where the capability of what's now known as the negative trust anchor was discussed and argued both ways inside of – it was first the DNSSEC Working Group, and then the DNSSEC Working Group at the IETF.

So the conclusion from the developers of the spec and the IETF published document is that there should not be a thing like the negative trust anchor, but as we've seen real roll out, you've got to do something to make sure your customers can continue to do what they want to do.

Anymore questions of the panel? Yes, Wes?

[WES HARDAKER]: Wes Hardaker [inaudible] Parsons. About half the room got that joke.

So you guys have had a lot of experience both as signers and as people overseeing customers that have tried to do DNSSEC or have done DNSSEC, and one of the fundamentally challenging things I think is that you have a different set of errors that you see when you're starting out versus the errors that you see long-running.

And if anybody – or all of you – want to comment on how can we help the people that they got over the first hurdle, what do we make sure that they need to realize they need to do next? What sort of errors are you seeing that it's like everybody's doing the same thing year out that they're making this mistake in the future once they get up and running?

UNIDENTIFIED MALE: I'll take first crack. There are a whole series of things with [inaudible] on automation and signing and validating that your software is correctly configured and various other things in choosing key expiration times when you first roll out and figuring out the first time usually how to get the DS up [into] the parent is someone problematic. Then it just sort of runs.

Actually, I worked for one of those appliance vendors in a previous job that the government was diving all over, and in most cases with the ZSK, they just had automatic rollover and that's what they like. You check this little tick box that says "Is this zone DNSSEC signed?" You check it once, and forever thereafter, you never touch it again.

The problem is that eventually someone comes along and says, "Best practice says you should roll over your key signing key every year," or every two years, every five years. And it's now been forever. The people

who have done it have possibly left and gone on to other projects, and someone either doesn't do that and the key signing key expires or you blow a key rollover. That tends to be more of where I see it.

The people who don't use an automated solution will occasionally [blow] zone signing keys. But usually it is a mismatch or someone mistakenly will say, "Oh, I don't want to do DNSSEC anymore. I will un-sign my zone, and don't take the DS record out," or the flipside.

DAN YORK: Jacques, did you want to respond, too, please?

JAQUES LATOUR: [inaudible] answer for [Sebastian] separate. So [inaudible] 93 domains. It comes from twelve different registrars.

UNIDENTIFIED MALE: That's interesting.

DAN YORK: That's actually a nice spread of registrars. Do you know what your total number of registrars is that you have? Is it in the hundreds or thousands?

UNIDENTIFIED MALE: 150-ish.

DAN YORK: So 10% of them are…?

UNIDENTIFIED MALE: Yeah.

DAN YORK: Okay, thank you.

[SEBASTIAN]: If you want, we can actually compare [inaudible] statistics here available, so we can [inaudible]. There was anyone waiting for questions? Because I wanted to [inaudible] to this discussion from Jeff and Wes.

So I don't know if you're aware or do you remember the SIDN, the .NL registry, they have an agreement with some of the ISPs doing validation for getting some kind of back channel of DNSSEC failure [inaudible] that go into the system and they fit that into the registrars. Because in the question of where the responsibility is, I think is all related to was the policy in the particular registry?

In our registry, [inaudible] registry, we have a policy that separates the registrars from the DNS operators.

So you might see in some cases the registrar is a DNS operator and then [inaudible] signing, but in some other cases, you'll see that's not the case. That responsibility is actually a little bit [inaudible] and hard to track.

And the other thing is you [can't] have your checks when the signed zone is out, but even like that, you might find some [traits] there or failures that are not detected by those checks, but you'll see them live.

It actually happened to us when we did the first signing [inaudible] zone that our [tool] had an encoding error that was only detected by [Nominom]. So [unbound] didn't see it. BIND didn't see it. We did all our testing and even like that, we got a call from Comcast saying, "You know your zone is broken." And I said, "What? How?" And that's the [inaudible] event.

So even like that, you need some live checking to make sure this zone is out there and it's being validated by different implementations.

RUSS MUNDY: Okay. One more quick, then we're about out of time.

UNIDENTIFIED MALE: Sorry. Not to beat the statistical analysis to death, but Duane – and I apologize if I missed this, but [inaudible] got a question for you. Of those four or five routers that you are highlighting, what was the percentage of your 400,000 zones that they accounted for? Did you cover that?

DUANE WESSELS: No, I didn't calculate that. The closest thing would be – the four or five that you're referring to, I think those are on the graphs at the end of the slides. I didn't calculate it specifically for them, but the very first graph has a distribution of the top signed registrars. If I had presented that a

little bit differently as maybe a cumulative distribution function, I think it would answer your question directly.

RUSS MUNDY: Thank you, everybody, in the audience for all these great questions, and especially thank you for our panelists for coming and doing these presentations. Thank you very much.

[applause]

DAN YORK: Okay. And [Gu] I need you to come up here, if you'd join us up here. I think I have everybody – I need Joe Abley, too, who is not yet here and Jim Galvin. There he is, okay. All right. And I need Joe who is in the [interop] room probably. He's coming.

So to those who are remote, as we are just settling in here, this is to be the panel discussion around the impact of root key rollover. And I think we're almost gathered with our people here.

So I'm Dan York. We're going to be talking today about what the potential impact is of the root key rollover and the pieces that are there. I want to begin with a quick round of introductions. And here's Joe. And we're putting him right there. Perfect. You get center stage. Right in the middle, Joe. Last one in gets to be right in the middle seat.

We've got a discussion here and this one is one that we're not going to sit here and run through some slides. We want to have some conversation around what this is about. We want to open it to you all who are here in the room. Again, as well, too, a reminder that this is a

bit of a preview of what will also be talked about tomorrow from 9:00 to 12:00 in the Brentwood room which is nearby here.

So let me begin by just asking each of the panelists to say briefly who they are and their connection in some way to what's going on with the KSK rollover. And then we'll get into more of what that is. Russ, go ahead.

RUSS MUNDY: I am the person from the SSAC which is, for those of you that aren't familiar, it's Security and Stability Advisory Committee of ICANN. We're the link of this workshop into ICANN, per se. And one of the things that SSAC does is publish documents with advice and recommendations.

And about a year ago, we published such a document on the root key rollover, and so I'll be talking about that mostly from that perspective.

DAN YORK: Joe?

JOE ABLEY: Hi, I'm Joe Abley. I was part of the team in 2010 that deployed DNSSEC in the root zone, and I've since escaped from ICANN. David Conrad waited until the root zone was signed before he left, but now he's come back. As an act of simple revenge, he's wrote me back into this exercise for rolling the root zone KSK.

You saw the large number of people just wandering in here. We've been across the corridor, a collection of vendors and DNS service operators

having a sort of pre-conversation about where we think the gaps are and what we think will explode and how brightly will the fire burn at various stages of this proposed key roll.

As was mentioned before, as Dan mentioned, we have also been roped into this more public conversation tomorrow about the wider aspects of it and how things should work. And I work at Dyn now. That's it.

PAUL EBERSMAN: Paul Ebersman. I work with Comcast, and actually our relationship to this is very simple. I am at the [wrong] end of 22 million phone calls if we don't roll this correctly.

JIM GALVIN: My name is Jim Galvin. I'm with Afilias. We are in this similar to Comcast, although we don't have quite as many users as 22 million. But we do have some 24 million domains under management. Not all signed, but we're ever hopeful that they're all going to be signed.

But in that context, as a critical infrastructure provider, we certainly care a great deal about how the root is managed. And in fact, as a host, a service provider for almost, when this gTLD program comes together, we'll have somewhere up towards 300 TLDs as part of our operation. So we care a great deal about what's in the root zone and how it's managed since it affects all of our customers as well as ourselves. Thank you.

YUNHONG GU: Hi, my name is Yunhong Gu. I work on Google public DS. If you've never heard about it, it's a DS server running on 8.8.8.8. It's public, open. So we do DNSSEC by default. We [inaudible] the DNSSEC response, about a few million responses per second. So it's pretty big. Thank you.

DAN YORK: Just a few validations at any given moment, right? So let me ask a question before we begin. With the folks who are here, how many of you are aware of what's going on with the root KSK? Okay, some. A good number. All right, good.

So we want to begin with just a brief introduction for those who may not know, and as well for those who are watching remotely about why are we here, why do we care about the root KSK, why does it matter? Jim Galvin has offered to do that.

I would also just say I've been asked to remind the panelists when we speak, when you answer a question or anything else, if you could state your name – and this goes for everybody else, too, if you're saying a question. The reason is because we also have transcribers who are creating a transcript and such, so they need to know who is speaking and they can't identify our voices without our names. And this is Dan York, by the way.

JIM GALVIN: And this will be Jim Galvin. Thank you, Dan. I offered just to give a quick overview of why this root key rollover is important and what it all means.

In simple terms, I think we can all appreciate that a key rollover seems like a very straightforward process. You create a new key pair, you distribute the new key pair, you sign with the new key pair and then you withdraw the old key that you're using. All of that seems pretty straightforward, and in general works very well.

The beauty of DNSSEC, especially with the DNS, is the DNS distributes all the key information and so it all just sort of takes care of itself. Aside from the fact that there's some timing considerations and implementation details. But the important characteristic there is it really does all take care of itself.

With that in mind, why do we really care here about this one? Well, the problem is we're talking about the root key and the root key is special. And the root key is special because it's the – being the top of the chain of authority, it's the one key that everybody has to know [inaudible]. The rest of the system takes care of itself, because as you step down through the chain of trust, you have a parent key which tells you that the next key down is the right one as you're making your way through the system.

But the root key, you have to start from somewhere and you start by having had it distributed. Everybody has to know it. Everybody having to know it includes resolvers – have to know in advance which one is the root key. Operating systems have to know in advance. Applications like browsers that have their own built-in resolvers have to know in advance.

So the problem is when I want to change the root key, I can't simply just put it in the DNS and let it take care of itself like it works with the rest of

the system. I actually have to tell everybody everywhere about the root key and it somehow has to get there and they have to know that they got the right one. And that's what makes this very complicated and that's an issue that we're here to be concerned about.

How do we set up a system in which everyone everywhere can get this key and have it happen in a relatively coordination, and hopefully mostly automated way without totally bringing down the entire system? Thanks.

DAN YORK: Thanks, Jim. Some of the other pieces that are in there, you'll hear a number of things thrown around here today like RFC 5011, which is an automated mechanism of updating the root key. You'll hear about manual root keys. You'll hear about manual trust anchor updates, other different pieces around that. We're going to talk about that.

Some of what happened last year was that ICANN had a public consultation that was opened up that was back in the spring of last year. A good number of people submitted comments from a wide range of organizations, a number of whom I can see in the room today who are here.

And out of that, the Security and Stability Advisory Committee (SSAC) created a report, and that report came out last year, which was then given to the Board to provide further updates on.

So Russ is going to give us a little bit of an update on what SSAC determined based on all of this input and their own research, etc.

RUSS MUNDY: Thank you, Dan. I was the person who was trying to herd all 40 SSAC members into a concurrence on a document about root key rollover, although many of the people on SSAC have been involved with DNSSEC, from the very beginning earlier stages. It was quite a challenge to get a conclusion that people would agree to.

And out of that effort the SSAC report actually contains five recommendations. Many of them are really tied to publicity and communication. As Dan has mentioned, as Jim has mentioned, when you change the root key, it literally impacts – it literally can spread anywhere in the Internet.

So we don't know where all the validation is occurring in the Internet, and validators in probably [99 in 100] are using the root key as their only trust anchor. Some little spots may focus on something different. But for all intents and purposes, the root key is the root key and it is what everyone is using.

So the really first and number one recommendation was get the word out. Communicate it out to people, so that they would in fact be aware that it was coming.

Now, the really big operators – in fact, I personally spoke with folks from Google and Comcast at the time. The people that are really aware of what's going on, it's no surprise to them. I mean, they watch these things. They understand it and they develop plans of what they're going to do, if you will. The biggies are not that much of a problem. It's beyond the large entities that pay close attention. Those are the places

where the problem can occur. And that's why that first recommendation was to really do effective communications.

I won't say any more about it in terms of how well it's being done. Same thing with the other four. But there is a webpage that you can go to to actually look at how the Board is keeping track of recommendations that have come to it from not just the SSAC but from other advisory committees, and you can see there the actions that part going on.

The next recommendation was that the ICANN staff needs to get active, get in the middle and do a collaborative effort to help with some kind of testing facilitation. Gee, I wonder what's going on today in this other room? So that's very much a direct tie to that next recommendation.

The third recommendation was to try to figure out what would actually constitute breakage. Okay, we're going to do this at some point in time. How do we know? How do we find out? How do we respond to something that's considered broken?

Next recommendation was to look at a plan and try to develop a plan in case "oh no, it went so wrong so bad we have to roll back to the previous key."

I don't know how much – actually, I'm not aware of much being done in that space, but that's something that the SSAC thought was quite important, being people that have seen things go much more wrong than they ever thought could have happened from their own personal experiences. That was where this recommendation came from. Be prepared, sort of a Boy Scout kind of approach.

And the last one was to try to develop in advance of doing this a set of information and data that you would want to collect and identify that could be useful for later subsequent KSK rollovers.

And I think it was Scott that was talking earlier in the enterprises in .gov how they'd start off good and then people would change contractors and the whole thing would fall back five years and they we remaking the same mistake they made five years ago or three years ago. The same sort of thing might happen here.

So the suggestion was get as much information as you can, identify in advance what you want to collect, figure out how you're going to collect it so it will be useful for the following key rollover.

DAN YORK: Thank you, Russ. That was a good summary of that report. I will mention, too, that if you go to the webpage for this DNSSEC workshop Russ has provided a set of slides that gives a bit of a summary of what the SSAC report is and you can also just search for that SSAC 063 and you should be able to find that report as well to see the full recommendations.

So let me ask a couple of questions here of the whole panel and maybe I'll begin with Joe. We're here. We've got specifications and such. I know you've been involved a little bit on what are some of the specification gaps, I guess I'd say, in terms of where do we need to look at some things that we've figured out now in the couple years since we've started to really work on what this look like that might need a bit more work in that space?

JOE ABLEY: I think we have a number of gaps in the way that the system was deployed in 2009-2010 that resulted almost directly from the fact that we didn't have a plan schedule for rolling the KSK and there's been no operational habit from rolling the KSK.

And these are not gaps – I'm not pointing the finger at myself and the team at ICANN. I'm pointing [inaudible] finger really at the whole system, because the lack of a key roll to date has really sort of made this kind of gap inevitable.

So I'm talking specifically about things like trust anchor retrieval and getting a secure copy out of [inaudible] at the DNS of the trust anchor that you should use and being able to authenticate it as being suitable for use for validation thereafter, assuming again that you have a trust anchor for the root zone, and that's how you manage your trust in the entire system.

I think it's probably fair to say that there's a large number of people who turned on DNSSEC because they were interested in it, and the way they retrieved the public key was to do dig.dnskey cut/paste. So they used a completely unauthenticated channel to retrieve it and they installed it. Nobody started to scream, so you assume that's probably fine then.

Trust at first use is not a horrible way of doing it. There's lots of examples where that's the most practical and sensible way of doing it, because the number of moving parts are very small.

However, there's no ability there to be able to manage the signaling that comes with 5011 with the publication of a new incoming key and the revocation of an old one that's an entirely manual process probably done by somebody who just turned it on for fun one day and then forgot about it, because as part of the standard 20-hour day you don't have a memory for everything that happens.

So there's probably a lot of unattended systems that are sitting around with a [hand] configured trust anchor that was not especially well-trusted, but I guess you can trust it now in the sense that if you still have customers, the DNS works.

And then correspondingly, as well as the initial sort of trust anchor retrieval, there's the more general bootstrap problem of what do you do when you do an upgrade of something or you replace a system and how do those things work?

As we imagine DNSSEC moving its way in through packages like DNSMasq and things like that into embedded devices and CPE, those things don't have administrators. So, worst case, we have hard coded trust anchors in devices that will never be upgraded and have no manager and we have no way of counting them.

We're heading back into a sort of situation – and I'll touch on the point that Russ made here about communication. We have the problem that we have a system that's critical for everybody who uses the Internet and we don't know who the users are individually and we don't know how to contact them. They don't know how to contact us in general.

And I'm not talking about the Comcasts and the Google Public DNS and people who are engaged. I'm talking about the long, long tail of people who have no idea about this process and turn this on because the [unbound] documentation said, "Just run this command and you'll be secure."

DAN YORK: Or even worse, it may have just been turned on for them. They bought the box and – boom – DNSSEC validation was just enabled out, which we're trying to encourage people to do.

JOE ABLEY: Right. So these kinds of gaps combined with these kinds of hands-off unmanaged system spell some kind of disaster – a disaster that we don't expect to be able to measure very well and we don't necessarily expect to hear from people about.

So this does not spell great confidence in how the system – it really means I think that the communication aspect is really, really critical.

So I'll draw a comparison. We have [inaudible] shirt on here. [Mark Sidon] made these shirts for SSAC [inaudible] controlled interruption. There we go.

So we have a situation with controlled interruption that was contentious, because it relied upon people who are lax enough in the administration of their systems that they leak internal-use TLDs to the root servers, but apparently diligent enough that they read their logs and notice an IP address which they then have to Google.

There was some doubt in certain people's minds as to whether this was a convincing way of providing a channel and finding out whether things broke.

But as it turns out, since this has started to happen, I heard the other day there have been 13 reports of people who reported seeing this thing and actually contacted ICANN to talk about – presumably, they were asked, "Was there loss of human life?" and they reassuringly said, "No, nobody died because of this." So that's okay then.

That seems like a ridiculous channel to expect any kind of feedback at all, and yet we got 13 responses.

When we rolled out DNSSEC in 2010, we had this address called root sign [inaudible] .org, which we published in the fake [inaudible] key for people who were looking there. Maybe that's equivalent level of expectation with looking in logs for name servers.

We did a massive public outreach exercise that far exceeds anything that happened for controlled interruption, and it was a much longer period for people to give feedback before, during and after the signing of the root. And how many e-mails did we ever get to root sign at icann.org? Any guesses?

UNIDENTIFIED MALE: Zero.

JOE ABLEY: Zero. So this communication piece – maybe that means that everything is fine and we don't have to worry. Or maybe it means that this is a

hopeless way of trying to contact people who don't understand the issues. I think communication is the central piece here.

JIM GALVIN: Just a little bit of a counter point to communication, though. And don't misunderstand, because I'm 100% behind communication. It absolutely is critical. But I think what's a useful observation here is the fact that some people don't even realize they're being communicating to, and that really is the problem that you're dealing with. It goes back to all the un-managed and [under] managed systems. I have no idea that this message is for me and I'm supposed to do something about it.

[DAN YORK]: Yeah. I think it will be interesting as we see all these devices that are out there doing it. Paul, you had mentioned earlier you're on the wrong end of 22 million customers that might suddenly lose access to – well, as Scott was saying, if 82% of the .gov sites are now being signed and something goes wrong with the root key rollover and suddenly all of you 18 million customers or whatever can't get to all the .gov sites – or 22 million, whatever that number was – it could have a whole lot of phone calls.

So what are you folks doing internally or thinking about with regard to this? How are you starting to think about this, or are you?

UNIDENTIFIED MALE: Well, the hopeful side of it is that we have automated config pushes and we have established procedures for doing any form of config change to

our DNS resolvers. And currently, everything that's in production or [exposed] on the Internet is in those processes. So it's one of those – "should" is a really ugly word, but it should be just a matter of simply doing config change push controlled across all of our various pops.

I am hopeful that whoever is doing this picks a really long tail for having the key overlap, however. One of the observations I'll make is we've had I think three root server IP address changes over the last two years, something like that. And the amount of time it takes to not have queries to the old IP addresses, even if it's just you hit it initially until you get your cache and load it, the tail has been horrendous and it instantly fixes, because as soon as you get on, you do a query and you actually cache it with the current stuff. So it's self-fixing.

The scary thing here is that the root key change is going to have just as long a tail and it will not self-fix. It will self-implode.

Fortunately, it's not my group, but I suspect we'll be dragged in. The place that I am concerned is we have a lot of business customers who are running their own resolvers and are not using ours. I suspect that our enterprise technical support team is going to be going through. So that's probably where I'll be putting the effort in terms of any form of training or [inaudible].

DAN YORK: That was a question I was actually just about to ask you, because for most of your residential customers, they're going to hit your main boxes. And looking at Gu down there, too, you guys have your main servers and if you update those, then life is good.

But for all those other folks who are using their own on your environment on the resolvers or are using their own and then chaining up to you in some way, that's probably your biggest exposure, right?

UNIDENTIFIED MALE: Yes. And based on things like open resolvers and broken CPEs and malware and various other things that we've seen, scarily enough, small- to medium-business is less diligent than home in keeping their networks clean.

So my expectation is that those will be the folks that have deployed things and forgotten or had somebody's nephew come in as a consultant and build things and have long since gone on.

DAN YORK: Gu, from Google's perspective, you're obviously a large provider of DNSSEC validation, what are you folks looking at in terms of how this all affects you or what you're doing?

YUNHONG GU: From Google Public's DNS [inaudible] about that. So we prepare for the [inaudible] somehow the root [case is wrong] or something happens. So we have a [inaudible] system where we monitor the root key, so in case it doesn't match whatever we have in the system, we just [inaudible].

DAN YORK: So you've set up your own monitoring system around that.

YUNHONG GU: Well [inaudible] resolver in the [world] [inaudible] don't have this monitoring system [inaudible] manage it. I think if we rollover a new KS, it's better to be in a [inaudible] manner [inaudible] we do all this to a small group [a region] or something like that, so we can monitor if something goes wrong. We can gradually expand [inaudible].

DAN YORK: So let's open up. I'd encourage people to start thinking of questions. I want to ask another one for the panel, but then I would really welcome questions from the group as well. You've got a group of people here with some good thinking around that.

I guess the question I'd ask to all of you guys is what do we do to make sure we do this right? Obviously we have the SSAC recommendations, as far as that plan. Is that really the plan that we should work on? Is it really [this] communication? Do we need test beds? How do we make sure that this doesn't blow up? Anyone want to…? Or do we all just agree on the SSAC report and go home? Jim wants to say something.

JIM GALVIN: What worries me most about the root key rollover is to build on something that Joe said and highlight a critical point. It's good to be talking about communication, but what we need for this to work right is automated distribution of the key. That is absolutely essential.

We have Comcast, we have Google and other large service providers. The benefit here is there are people involved in the process. Ultimately they'll solve the people side of this and the distribution side for the people thing.

But as Joe made reference to, the problem is what do you do about the unmanaged objects that are out there, the unmanaged things where there's no human intervention that's possible?

What's missing here is all of the applications, the service providers, your operating systems, the web browsers, they need a mechanism just like Comcast and Google have for managing the people side of this for the clients that have no idea. And being able to update those things and do that in an emergency way and getting that done.

DAN YORK: So do we, as a community, need to be engaging and interacting with those application vendors, with the [CP] providers? Where do you see us needing to talk to?

JIM GALVING: I would say yes. I think a part of the communication channel – when we talk about the communication channel, we often talk about how do we tell people about the new key and make sure they have the right key?

But I think part of the communication channel is the development of the importance and significance of the root key rollover, of security in general. We're so fond in general in the security space of saying security is an afterthought. People don't design it in in the first place.

This is a case where designing it in in the first place is absolutely essential. We're well along in the process of not having designed it in, but the next generation, it can't go there. It can't go there. This is the critical thing – designing it from the start.

So the communication part has to be about getting everybody to the table – and I really do mean everybody – to think about how they're going to be updating their systems when they need a new root key.

DAN YORK: Russ?

RUSS MUNDY: Yeah. Obviously I think the SSAC plan is a good idea, but the SSAC plan is incredibly abstract. It really doesn't say anything about the reality of what you do to execute it. And frankly, that's the way SSAC usually talks. SSAC is there to give advice and guidance, not to tell people exactly how to do things.

In fact, there's a whole bunch of the "how to do things" that are buried inside of each one of these recommendations, but of particular importance is to have the critical technical pieces in place and actually try them out before you do it for real. That's one of the ones that underneath of a test bed, doing some testing, getting outreach, as Jim said, to everybody you can get a hold of.

DAN YORK: So from a technical point of view then, let's just maybe talk about that for a second, what do we need as a community? What do we need from

a technical point of view to be able to test this? Or is that what was being discussed in the other room over there? I'm looking Joe.

JOE ABLEY: The flip answer is we need Geoff Huston.

[laughter]

I can offer you some perspectives on precisely the size and the issue of the problem you're facing. We have done extensive in looking at the resolvers—

GEOFF HUSTON: This is Geoff Huston, by the way, for the transcript.

[JOE ABLEY]: Oh, so I'm [inaudible] APNIC. We've looked at the resolvers people use. The [inaudible] stat is, as I've got here, 75% of [inaudible] have their queries forwarded via just 1% of visible resolvers.

So if you take the top 33,000 resolvers – and that includes the Comcasts and Googles and open DNSs of this world – but if you take that 1% of resolvers, you've got 75% of end users covered already.

We actually see 25% of users do DNS key queries from resolvers. It's really hard to figure out which resolvers are validating. Just have to take the DNS is a chaotic mess and you just can't see inside. There's no X-ray vision.

But there is hope that, quite frankly, the number of DNSSEC validating resolvers out there is identifiable. In other words, even today we know where the problem is for my suspicion around 95 to 99% of all end users. We know who will be affected by root key rollover, because again, what I see is that trailing – and it looks to be around the trailing, 90% of resolvers are used by one or two clients that don't do DNSSEC validation. So most of those 200,000 or 300,000 resolvers aren't your problem today.

So there is hope and we do need some better data. And I think with a bit more time and digging, I can quantify this and give you long lists. But yes, we know kind of where you need to look for the resolvers that are going to be impacted.

DAN YORK: So to your point – okay, so if 1% of the resolvers are doing 75% of the queries out there, then what we—

[JOE ABLEY]: No, 75% of the end users.

DAN YORK: Of the end users, okay. All right. So what you're basically saying is we need to contact that 1% and say, "Hey, guys, make sure you're doing this," and then we're done, right, for 75% of the users?

[JOE ABLEY]: Pretty much.

DAN YORK: Okay, all right. This is a new meaning of the 1% or something like that. I see Warren, yes?

WARREN KUMARI: Warren Kumari, Google and also SSAC person who contributed to Russ's report. If we assume that the top few resolvers are well-managed and the rest of them aren't nearly as closely managed, we can hopefully assume that the top X will get fixed when this happens.

However, as time goes by, more and more and more of the rest of the resolvers will be turning on validation, we assume. This means that I think we need to try and do this as soon as possible. As more unsophisticated users start doing DNSSEC, there's likely to be more breakage. So if a certain percent are going to break, that percentage or that number of users affected will be smaller now than it will tomorrow and the next day and the day after that. So we should just do this.

DAN YORK: Sure. There's some agreement around here that I think we're looking at. As a guy who my job is partly DNSSEC advocacy, every day that I'm talking more about this and encouraging more people to do it means more people who could potentially break when this happens. So I'd like this to happen – tomorrow would be awesome. But I don't think that's going to quite go.

Joe wants to say something.

JOE ABLEY: Yes. All the stuff you said is fantastic. It all makes sense. I think what we have to do is keep our eye on what the objective is here. All the comments we've had so far, possibly apart from Jeff's, have been unbelievably pessimistic. If you are not used to this kind of crowd and used to this kind of discussion, you would imagine that the only natural conclusion is, yeah, we should not do this.

In fact, I think the desired outcome is the opposite. I think not only do we want to do this and shake out the bugs in the system at this point, but we want to get into the operational habit of expecting this so that you routinely have [code] that make bad assumptions about [hard-coded] trust anchors fail. And not just once every five years or once every 15 years, but on a much more rapid schedule – a much higher frequency – so that this stuff comes out.

The end goal here I think is that we have a system which not only features regular and reasonably-high frequency key rollover, but also that we learn from those scheduled key rolls and become better able to deal with an emergency roll. The ability we have at the moment to deal with a key compromise, for example, is a complete trust [inaudible], with no notice and no communication which is going to cause far more disruption than the incremental pain that we feel within the long tail of people who are not engaged of doing the structured organized key roll.

So despite all the pessimism, which is natural with engineers because there's always this great resistance to any plan that's going to cause a phone call at 3:00 AM. The goal here is actually to endure that pain so that, in the longer term, we can actually sleep at night.

DAN YORK: Yeah. And let's be realistic. We are in an even a greater subset of the engineers, which is the security engineers [in the] people, and we just have this perpetual – we're paranoid downers on things as far as we see the worst in all of that. So, yes, we have to remember about what's going on there.

I have a queue of people who want to jump in here. Warren, did you want to just quickly respond to that? No, okay. So, I've got Jim, Mehmet and Russ – no, somebody else was in there. Paul, sorry. Mehmet – no, Jim. Sorry.

JIM GALVIN: As I listen to Jeff's statistics, which are very, very optimistic and positive – that is such good new. If I get 1% of the resolvers, I can take care of 75% of the end user community perhaps.

But what concerns me with focusing on that is coming back to what Joe said. What is the problem that we're trying to solve?

I think that, while that's good news what Jeff said as far as the statistics are concerned, that means that we've got some hope of getting pretty good, we really need for the root key to roll regular, to get good at rolling it, to expect it to change and it has to happen everywhere at the same time all the time. You're not allowed to have a long tail. The problem we're trying to solve is to not have a long tail of [inaudible].

DAN YORK: Which is great. I think we can all agree that we want to – well, we may not all agree in this room. But for the purpose of today, let's focus on

the fact that we do need to roll on this. I really want to talk about how we roll. What are the issues that we've dealt with here as far as that?

So, Mehmet, I see you.

MEHMET AKCIN: I had exactly the same points, actually. Both with Joe and Jim. So I'm not going to add into that. But just one quick thing.

I have been reading the SSAC report on this. I'm kind of surprised, actually, to not see a recommendation towards regular rollovers. I think it should be there. It should be recommended as, "Hey, this should be regular practice."

I'm not saying specify a timeframe, because that's not really what SSAC does. But I think it should be there as a recommendation.

DAN YORK: Okay. The SSAC provided the advice to the Board, which then voted on that last November to ask the ICANN staff to prepare recommendations that met with this advice. So that effort is still underway to provide a plan to implement the recommendations, would be the precise way to say that. And that's part of what this workshop yesterday was, what the workshop today is, and part of what the workshop tomorrow is about, too.

RUSS MUNDY: I'd like to just thank Mehmet for making that suggestion and let you know that putting an SSAC report together is a lot like making sausage.

You never know what's going to come out the end until you see the end. And in the midst of the sausage production, there certainly was discussions about that suggestion for recommending that it be done regularly and it didn't make it into the sausage.

DAN YORK: And by the way, while I'm here, I will just say, Mehmet, thank you. We did thank the sponsors this morning for the workshop. Mehmet is with Microsoft and is responsible for the sponsorship that they had here. So thank you while you're here.

Moving on to Paul.

PAUL EBERSMAN: I think one of the things that we can do to move towards automation, which I think is going to be the only long-term survivable is much in the same way that the 1% hits 75% of the users. There's probably a fairly small number of places where people are getting resolvers that are their own. There are companies that are running internal resolvers, which are probably the big three, and they're probably using [distrose].

So we go to the [devians] in those and say, "Put this into your next apt-get version. Microsoft corporate, anybody who has a support contract is already getting automatic updates.

The other big piece, the bane of our existence in many ways with open resolvers are CPEs, many of whom are using DNSMasq because it was free and small and lightweight.

It has been updated recently to do DNSSEC. CPE vendors are very slow to roll things out, but that's going to be another one where before they start rolling [inaudible] need another version, perhaps we should be talking about a way to automate that so that home routers that got kicked in also pick up the new keys by whatever automation we decide [to use].

DAN YORK: Sounds good. And I would mention that in our very next panel we're going to have our operating systems, so we'll be able to hear from Microsoft, FreeBSD, and Fedora. So we can ask that question in there, too.

I see Geoff.

GEOFF HUSTON: Geoff Huston, APNIC. There are three questions in my mind about the root key rollover. The first one I've talked about, which is the resolver population. Who is actually affected? There are two other questions as far as I'm aware, and there are probably a lot more [than] those two.

The first is support for RFC 5011, old key signs new. As far as I can tell, I cannot test our system today and give end users who don't know what's going on a simple DNS test that will expose their capability. I can't do that.

However, the third question is equally interesting, which is when you make the response size bigger as you do in a 5011 style rollover where

you're sending back more keying information, there are two big barrier points that we actually observe out there – 1280 and 1500.

What happens to users when the DNS response side bundled up goes past those two thresholds? I can tell you. I can measure that. We've done some preliminary measurements last year, but it is a kind of tricky one, and certainly our next set of measurements that we're going to do is actually isolating the dropout rate, identifying which customers behind which networks actually don't get it, because the technique is good enough to actually isolate in where those problems lie.

So those are the three that I see. If there are other questions that are of interest that we can measure around the existing DNS system, I'd be happy to have a look at them. Thanks.

I see Warren right next to him.

WARREN KUMARI: So let's say we go off and test and we discover that a bunch of networks don't work correctly with 5011 or with some part of this new DNSSEC key roll thing. What do we do at that point? Is somebody going to take on outreach to specific networks and say, "By the way, this is coming. It's not going to work." Is that something potentially ICANN could do as well as a general outreach to let people know that key rollover is happening. Perhaps they should be targeted and go find the following [end networks] and poke them, please.

DAN YORK: I would say it's probably a combination of both of those things, as if we can measure it to a certain degree. Then that comes up with the possibility of if we know where the brokenness is, then we can start to try to identify some of those folks.

And maybe it's a website that has the statistics, and then you start to do the general promotion, get it out there to the operator group mailing list, get it out there to the other different places and try to do what you can in a public relations kind of form to go and promote that.

It certainly is a thing that I think ICANN and the broader community of all of us in this room can be involved with that kind of effort. I'm looking at Geoff as far as the one with the metrics. People like Geoff who can create that are certainly helpful in that.

GEOFF HUSTON: I probably should – or at least I gave a slide presentation at DNS-OARC on Sunday and I listed 30 networks whose resolvers did not support [ECDSA] for the customers. I believe three of those have been fixed already.

Now, oddly enough, naming and shaming in this community, that's not really shaming – but naming and identifying where there are problems, particularly on the DNSSEC community, it elicits a response saying, "Can we fix this and fix it now?"

So I suspect a combination of the two. If the lists are easily identifiable, certainly in terms of a root key rollover, there is a responsive bunch of folks going, "If there's something we can fix, we will fix it." But I'm

pretty impressed that within three days, three rather large networks have now got ECDSA support.

DAN YORK: That's pretty cool. Actually, I'll put a plug in that tomorrow at the public workshop on the KSK, Geoff is going to be giving part of that presentation – or do you know that? All right, good. Just checking.

Geoff is going to be giving a bit of what he gave at DNS-OARC last week about a potential thing with [inaudible], which would basically be what if we changed the algorithm, in addition to rolling the key for just getting a new key, there's other aspects of this which are things like what if we want to use a new algorithm instead of the existing algorithm? What if we want to use a longer key rather than a shorter key? That type of thing. So there's additional issues, rather than just we want to put in a new key. Reasons for that. Another tip for tomorrow. Oh, you have something?

RUSS MUNDY: I wanted to respond to Geoff's other two questions with a little bit more background material.

There was a 5011 test bed running about three to five years ago as part of the DNSSEC deployment project. It was shocking how well it worked. Everything just worked. Because everything was working so well at that time, it basically was shut down. And I know there's efforts to get another one and do it up and that's a very good thing. I think it should happen to be ongoing.

The other aspect, the large response size, we do have an appendix – Appendix B in SAC63 – that has a bunch of details and identifies things to look out for in that space. So I commend folks that are interested in some technical details. Go take a look at Appendix B in SAC63.

DAN YORK: I've got Joe and Jim in queue and I've got 15 minutes left. If there are people in here – oh, and Jacques wants to be in here, too. Warren, too. And Mehmet. All right, good. If there's other people behind me, too, let me know as well. We'd like to incorporate as many voices in here.

Geoff, I would also agree with you. One of the things that's interesting with those DNSSEC deployment maps is, because we publish those and they go out every week, I've had inquiries from various ccTLDs who say, "Hmm, look. All the people around us are signed and we're not. Maybe we should do something about that." So there is. It's not shaming, but it's just shining light on the people who are doing the good things and providing some incentive to that.

Okay. I've got Joe and then Jim.

JOE ABLEY: On the general topic of measurement, Geoff, we've heard many words from Geoff today as we usually enjoy on how we do real user monitoring. That was not intended to be sarcastic. Where was I?

I was going to think back to 2010 when we did DNSSEC deployment and we deliberately rolled out the [inaudible] server by server or in groups of servers, so that at any one time during that window we had some

servers that gave the normal response sizes for referrals and we had others that gave inflated responses with signatures for clients that had the [DO bit set] regardless of the fact you couldn't validate.

We collected a lot of information. I think the longest period [inaudible] that OARC has ever done. Duane over here spent a lot of time mining through that data to try and find what was going on.

My point really is that we have two sides that we can measure. We can measure the real user experience in terms of what works, what doesn't, what gets responses, what doesn't. And we also have the ability, which historically has been a fairly cumbersome sort of effort to try and instrument the root servers.

The thing that's different this time around potentially is that the new improved with additional VIM RSSAC, which is poised to release its first two documents. The second of those documents (number 002) is a measurement document, which describes a framework of precisely the kinds of metrics that root servers should collect and to change format, and then ultimately – hopefully – a central repository where there's data can be deposited and something resembling real time. Not just for a particular schedule DITL event, but just in general.

So, potentially, we're in much better shape to be able to measure this kind of stuff as it goes on.

The other thing I wanted to mention that connected with the idea of the DURZ and rolling this thing out piecemeal across root servers is we have the beginnings of an idea that Roy has recently added a twist to, to say that while we roll the key and we have different root zones with

different sets of [DL] signatures and additional KSKs published in the DNS key [RO] set, there's also no reason why we have to have that coherent across all the root servers.

So as that goes out, and particularly around the trigger points of things like the revocation bit being set in the outgoing key or the DNS Key response size increasing or the first [DL] signature that we have by both KSKs, we have the ability potentially to be able to watch validators and resolver traffic swing from root server to root server. It's something more resembling real time than we could do in 2010.

So I think we can approach this from both sides and I think the measurement picture is actually much better this time than it was back four years ago.

DAN YORK: Cool. Jim?

JIM GALVIN: Thank you. I want to come back, Dan, to your original question here about what can we do to make all this work. I wanted to response directly to one of them. So I think I'm hearing three things.

We're talking a lot about metrics and measurements. We're a room full of technical people. That's where we tend to go. That's sort of the foundation. We need to see what the state is now and we need to be able to see how we're doing as we go along.

The other suggestion we have is about automation. I think we all agree that we need automation. What I want to add to that is I think there

needs to be layers of automation and we're going to build this up as we go, as we measure. We're going to be finding these edge cases and looking for better ways to improve that. ICANN has to announce a new key. There's a whole set of service providers that have to build up mechanisms for grabbing that key. They have to have their own distribution mechanisms and all of that has to happen. We're going to find the edge cases and the people who are not at the table.

Communication, though, is really the hard one. Warren asked a question about who does the communication? Who's responsible for it?

What I want to say there is, just like with automation, there are going to be layers of communication that have to happen. In fact, just like the key has to be everywhere, everybody has to be part of that communication and the communication is going to grow. We have to roll the key, we have to do it now and we have to roll it often and continue to test as we go.

I say those. I don't mean those as a deliberate go do it tomorrow and the next day and every hour or something. Just as a principle, we need to begin to roll it and to roll it frequently and regularly and make this a habit and get into it.

More people are going to become involved in communicating. As Geoff said, it wasn't a wall of shame when he put up the list, but as more people talk about, they see it, they're going to want to fix it.

What I worry about in communication – and I think this is important, that they'll be these layers. What I worry about in communicate coming back to end with one point that I made before, I worry about the people

who don't know that the message is for them. That's what I worry about. I don't know that we solve that problem, except by continuing to do it and continuing to grow our information and continuing to tell more people. Maybe that problem fixes itself.

But that's the problem I see in communication that's difficult to solve – the people who don't know the message is for them.

DAN YORK: It suddenly pops into my mind, I know the Consumer Electronics Association (CEA) has an IPv6 working group, but I'm wondering – we should also be talking to people like them on the embedded device side around the DNS and DNSSEC side of things.

I have Jacques.

JACQUES LATOUR: Hi, Jacques with .CA. I guess I have a naïve question. I'm not a rocket scientist with all [inaudible] and stuff. But I'd like to understand what the thinking was behind having the trust anchor out of [bend] instead of having it inside the DNS. Because I don't see the trust. I don't see it being more trusted as an html file or an xml file somewhere than having it inside DNS. If it's in the DNS, then it can automate more stuff.

UNIDENTIFIED MALE: Rick Lamb, who is sitting behind me, can probably speak more eloquently to this. Do you want to talk to this or shall I just…? Okay. We don't have a lot of time to drill down into huge depth anyway.

But the idea was that we expected this trust anchor to need to be sent out to applications and to operating systems and that we understood that software distribution and software update cycles had a natural mechanism for doing this already. They had keys and the ability to validate signed code, so that if we could package the trust anchor in such a way that it could be validated and trusted to the same level as the code that's running in your application and your browser or on your laptop or whatever else like that, then we have a similar level of trust that's acceptable in order to bootstrap the entire thing.

So the idea of not doing it in the DNS was to try and fit as closely as possible into what we imagine the most normal use cases would be. As it turns out, as I mentioned before, there's not been significant interest or uptake that's been obvious for this since that happened. But that was the original thinking.

RUSS MUNDY: I'd like to just add a little bit. It really is also a result of the fundamental protocol design. It was designed from day one so that the validator side of things would likely use the top-level trust anchor, but there was also a perceived requirement. And honestly, I don't know how much it's being used, for validators to be able to insert their own trust anchors. That could be from wherever.

So the hierarchical structure was the thinking to begin with and also the addition of individual trust anchors that users might want to make use of.

There was also a workshop – I believe it was in 2009 – that I can probably dig up the report from. It was a separate workshop to talk about now that the root is signed, or about to be signed, what next? We hit on a bunch of these questions. So if there's interest in that, I'll dig up that report for folks.

DAN YORK: Okay. Two last quick questions. Warren, you've got one in there? I don't know. You were in the queue.

WARREN KUMARI: Russ had mentioned something about root zone test infrastructure that we had for a while. We have a very, very rudimentary one if people want to test their 5011 at a site called keyroll.systems. So if you run 5011 and want to know if it works or if you implement the name server and want to know if your 5011 stuff works, try a poke at that.

DAN YORK: Is that actually a new gTLD in usage there?

WARREN KUMARI: Uh-huh! It was specifically registered there just to see if they work, and so far it has. If it doesn't work, keyroll.snausages.com.

[laughter]

DAN YORK: Mehmet?

MEHMET AKCIN: So I realize that DNSSEC key is created for DNS, but I would like to emphasize the fact that this is a certificate that was transparently generated, and this might also be embedded in some non-DNS related services.

I think when ICANN staff is looking to roll over this key, it should be also kept in mind that there might be people who have hardcoded this key on a non-DNS related thing. I'm sure Joe knows what I'm talking about here and Rick. They need to be aware and they need to be prepared about updating this as well, because they will need to update some things and we won't be able to see them in the resolvers because they're not using DNS for validation, but some other purposes.

KUMAR ASHUTOSH: I have a [inaudible] to that, actually.

DAN YORK: Okay.

KUMAR ASHUTOSH: This is Kumar Ashutosh from Microsoft, Windows DNS server. What I wanted to say, Mehmet mentioned one of the cases and there are many cases out there which we know, which we do not know. Things like root key rollover, which are going to impact like everyone.

Don't you think there should be a midway somewhere, a pilot where both the old key works for sometime and this thing [inaudible] rolling over so that people can verify, okay, this is going fine?

I know the RFC 5011 and other [inaudible] mechanisms that provide the bandwidth to get your trust anchors updated and stuff. But what I am saying is we do not really know what are the [inaudible] that we can black out people. Don't you think we should probably concentrate on finding out our test mechanism or a pilot mechanism to do such thing at such a large scale? Probably it's something that a charter IETF can [think of], and how to do this as a test pilot. That's all.

DAN YORK: Agreed. Joe wants to say something.

JOE ABLEY: I agree completely, and I think a fundamental component of all the draft plans that have been discussed to date for how, operationally, a KSK roll in the root zone would happen include a test bed – a public test bed – for both for people who run DNS services that need to test validation and their reaction to it, and also for people – vendors – who actually produce the software. So I agree.

KUMAR ASHUTOSH: I have a different opinion from the test bed part. I'm saying a test bed requires some kind of subscription where people come to you and then they test their systems. What I'm saying is in the intermediate we have a plan where people who are not even subscribing to your systems but

are going to use it, they find out that this is going on and they find out without blacking out their systems. It's a different approach.

They find out that this has gone on and they found out without blacking out the applications they are trying to connect. So some way of telling them, "Okay, this is not working in your case, but yes, it is going on there." It's very abstract what I'm thinking, but yes.

DAN YORK: Yeah. That kind of thinking is definitely welcome, because we'd like to understand what are mechanisms that we can do that. It occurs to me that there's been sudden efforts by the, for instance, the browser vendors who have put up different pages when you go to a site that is using old SSL or something like that, or different types of things saying you need an updated version of this.

There may be some mechanisms we can look at in that kind of space, which is a perfect segue into kind of the final minute or so of this to say, for some next steps, if you're interested – and, actually, that would be a great item to bring, for instance, to the discussion that's happening tomorrow from 9:00 to noon in the Brentwood room, which is somewhere right over here.

That is going to be a discussion that's going to involve a number of the people who are on this panel and other folks in this room around – in a three-hour time period really more of diving down into what do we do? What are the next levels of things? Where do we take this in terms of what are some of the larger issues? Geoff will be talking again about algorithm roll. We talked a bit about key size kind of issues and

identifying what are some of those pieces. And then who do we need to reach – some of the pieces around that – and how do we get there?

So that's going to be kind of the next step coming out of this is I would encourage people to come do that tomorrow. If you have people who are remote, that will be broadcast remotely so they can be able to participate in that.

There is also a mailing list called the KSK-rollover mailing list. I put up a post on the Deploy360 blog on the Internet Society site that has some of these links in there, which you can find out the link and how to go and subscribe to that, otherwise if you probably just search on KSK-rollover, you could find it I would imagine, or KSK-rollover mailing list or something like that.

But that list is having some good discussion about these kinds of issues, and I would encourage you to join that list if you're interested. I think if I could summarize, the general point is that this is going to be a massive communications effort, certainly, to go in and help people understand across all the industries.

It's not a one day we have the old key and tomorrow we have the new key. It is a staggered role. It is a period where there will be both keys in the current plans, so that it won't be something that just happens [inaudible]. But there will be a time when the old key is withdrawn. And at that point, that's when we expect things will start to break. So there will be a time period in there where people will be able to go and use both keys and do that.

With that, I would like to thank our panelists for their participation this morning.

[applause]

And now we have a break for 15 minutes. There should be food – or drinks? No, there's nothing. All right, it's a break! Stand up, stretch your muscles, come back here. And we need the operating system to come back here.

[break]

UNIDENTIFIED FEMALE: Please take your seats. Break is over. 45 minutes and you'll get another one for lunch. Ding-ding-ding! Break over.

RUSS MUNDY: All right, folks. Please come back in. We want to get underway. We're very close to on time and we want to remain that way. In fact, I will go close the door. Okay, session has begun. Thanks, everybody, for making it back. Hopefully we got a little bit of a chance to stand up and stretch our legs.

This is our panel on operating systems and DNSSEC implementation in operating systems. We have three operating systems represented here and the presenters are associated with the organization, except for one, and I'll take care of that when we get down to the Paul Wouters introduction. For those of you that might happen to know Paul, you won't see him on the [dais], so Wes is substituting for him.

Ashutosh is our presenter from Microsoft and has kindly agreed to tell us where Microsoft is at and where they are going with DNSSEC and Windows.

KUMAR ASHUTOSH: Hello, everyone. It's really a pleasure to be here. I'll just move to the presentation. Can you go to the next slide?

Just a brief overview. Windows DNS server is widely deployed in the enterprises, particularly. We are fairly present in the DNS resolver space. We are standards compliant and interoperable and secure and scalable. When I say secure, we especially insist on the DNSSEC implementation that we have. Next.

So this is pretty old. In 2008 R2 Windows DNS server actually introduced offline signing where you can sign the zone offline, create your [keys] offline and then put the zone online.

But next when we moved forward, DNSSEC Windows DNS server has come a long way. It is [inaudible] support. It a beautifully-implemented NSEC3 support. We support RSA/SHA-2 algorithm and ECDSA signing.

The trust anchor rollover is totally automated, 5011 is supported, and we also support the third-party key management apart from our own Microsoft [inaudible] service providers that you can bring any [CNG] compliant key management devices and that you can plug into Microsoft DNS servers and they will use those keys to sign. And if you want to implement [CNG], you can implement your own. Next.

So the major support that we have is for online zone signing, and when I say online zone signing, you can think of DDNS [inaudible] where you can just say dynamic updates and keep on adding records into that zone and it will automatically sign those new records without any delay, and it will also create the NSEC3 chain if the zone is signed NSEC3 parameters. And if they are removed or updated, the zone signing is updated with that.

We have worked very hard on improving the DNSSEC server performance, particularly the signing time, the validation time, and the response time because DNSSEC automatically comes with a lot of data, a larger throughput. So if your network has larger throughput, we at present, are [at par with] [inaudible] DNS server performance of Windows DNS server.

The second important part is the trust anchor management. We have a specific case for root trust anchor management. You can specify a URL and that URL will just download the root trust anchor, wherever it is.

We have [inaudible] specific trust anchors if you have, say, within the enterprise or if you want to have a particular trust anchor, you can import them as anything and these are automated very simply. You can import a file or you can just point to a DS set that you can import as trust anchor.

Then we have signed delegations for internal names, which we have Windows Parent and Windows Child. If you have a DNS server which hosts both the parent and child delegations, then what happens is it automatically updates the parent [inaudible], and when it is rolling over,

it will find out whether the parent DS has been updated or not and [inaudible] rollover, so that your whole chain of trust is not broken.

And 5011, yes, we do support. So from the root key rollover perspective or any trust case key rollover perspective, it's pretty fine. We are up there. Next.

Now, what I really want to insist here – and we at Microsoft, we stress the user benefits how easy it is to use. DNSSEC has been probably most complicated of the DNS operations stuff that has come out in a long time.

And what drives people away from DNSSEC is the complexity of the management. You have so many steps to do, so much things to configure, key generations stuff and that. Windows DNS server, it is simply a three-click process. You just click on a zone as shown in this picture. You click on the DNSSEC tab, you sign the zone. You have the option to go and explore all those variations. But there is a default signing parameter set which has been worked together with the U.S. government and the [inaudible] rules and this works pretty fine. It just signs every time, like all the recommended parameters.

There is a partial support, because many of the people here are not [inaudible], you scripting guys. We have Windows have the [inaudible] which is a very structured scripting language in which you can actually objectify the data that is coming out.

For example, if you have a full zone signing data, then it will be structured data rather than a plain flat file, where you can get that partial support in Windows DNS server.

Manageability is something that I really want. If you want to try out, we have people here. We can just show you how easy it is to manage the Windows DNS server.

And same on the resolver side. You have just got to do nothing. You just import a trust anchor in the simplest possible way. It [rolls].

The last thing that I want to emphasize on this is the automation. As I've said, automation is also an important part where you do not want to intervene. When a DNSSEC system is deployed, you do not want anyone to keep on managing things.

So whether the records are moving in and out, which really do happen when there are devices moving out or the application servers or cloud infrastructure which keeps on servicing [VIMs], the resigning of static and dynamic updates also happens.

Key rollovers are automated. Nobody needs to [inaudible]. You just specify the timings that work. Signature refreshes are automated. Secure delegations are automated and trust anchor updates are automated. Next.

So how is the signing [inaudible] process very simple. Generate keys if you have a CNG-compliant third-party KSPs [inaudible]. We introduced a new system called Key Master because I will not be able to cover everything, so just let me emphasize on that. We introduced a system called Key Master, because we predominately work with active directory systems.

So [an active directory], these are distributed multiple instances of the same zone – a replica of the same zone – on different servers. So you

sign on one zone, one server. The active directory will take care that the signature is replicated without your intervention to all other places.

Even if you have [inaudible] or non-active directory integrated zone, you can just write a simple partial script to automate that process that it will automatically do.

Your signatures – your KSKs in particular – are kept not in the [DMG] in a secure process. You can just use zone signing keys on the [DMGs] to sign the zone. Your keys will not be put out to the world. So that is how we also comply to the FISMA requirements that you should not expose your KSKs out, signing data outside. Next.

So Key Master Role I just talked. Can you go next? Next, next.

Zone signing key rollover, we subscribed to that 4641 mechanism – RFC 4641 – and pre-published mechanism for the [inaudible] and the double signature for KSK.

We have our own [inaudible] implementations for the trust anchor updates, and trust anchors, if you have added, they will automatically update from the source of the trust anchor.

And we have 5011 and [inaudible] time periods configured. If you want to revoke a key or retire a key, that also complies with the rollover process and it is all automated. That is what I want to insist.

These things, I had recently in a different session on DNS-OARC, I had mentioned that Microsoft has come with [inaudible] management policies in which you can do something very similar to CDNs which you

can use the DNS infrastructure to redirect your queries to intelligent responses.

If you have people coming from [inaudible] they can be redirected to the [inaudible] the people coming from Europe can redirected to the [inaudible], which was not possible using the classic DNS infrastructure.

And the biggest problem and that solution was having DNSSEC implementation. If you have intelligent responses which you are calculating online, then you require online signing, because DNSSEC signs the [inaudible] not the [inaudible].

We have found a beautiful way to support that, and intelligent responses are supported in policies. This is all I had to say. I think I'm well in time.

RUSS MUNDY: Well in time, indeed. Thank you. Very interesting presentation. It's really great to see the progress that Microsoft products have made with respect to DNSSEC, in particular the automation aspects. I think that's wonderful. Okay, thank you.

Our next presenter is Erwin Lansing of the FreeBSD project, and he will be telling us about where FreeBSD is at and where they are headed. Thank you, Erin.

ERWIN LANSING: Thanks, Russ. So I'm going to focus a bit more on the resolving side for the authoritative and signing parts. We got packages for BIND, not

[inaudible] OpenDNSSEC, you name it. You can just [inaudible] with that. First slide, please.

So for those of you who were in [inaudible] in the tech day, you might remember this slide. It looks almost similar, except at the time it says "work in progress" and now it's done. We imported unbound into the base system as a local caching validating resolver. Of course we also used ldns and those tools that come with that [inaudible] finally gone. [inaudible] is now [inaudible] and we found a host-wrapper, so host is still there and does whatever we did with the same flags as usual.

When I say a local caching resolver, it's meant only for local host. If you want to do your traditional site-wide implementation of unbound or any other resolver, install that from a third-party package instead and turn off the local resolver. It's really just meant for the local [machine].

It basically just takes whatever [inaudible] insert itself into the middle of that and forwards the queries to whatever [inaudible]. But just a validation in between no matter what your forwarder is doing.

Of course we also check for SSHFP and the whole project loses DANE records for all the certificates we have in our own [inaudible] zone. Next slide.

So that was 10.0 that came out a half-a-year-ago, a year ago. 10.1 is on its way right now. It should be finished next week, hopefully. There were some minor issues. We resolved all those [inaudible] and we still haven't turned it on by default, but it will now be a checkbox during the installation just for the visibility for people to see that it's there. Next slide.

And hopefully for 11.0, it will be default on. Next slide.

Some of the issues we encountered is things that turn up in our DNSSEC deployment is getting more wide and people start to see all kinds of use cases that do not work anymore.

One interesting thing is if your time is off too far where you boot your server, you cannot validate your signatures because your timing is wrong and you cannot check. You don't get your DNSSEC records right, which also means you cannot look up your NTP servers to get your time right. So now we have to call a human to just set the time manually.

The other thing which is even more interesting if you run around with your laptop at these kinds of conferences is what we call the "Starbucks Effect." We had a very interesting case back in May when we had our yearly conference in Ottawa where the University of Ottawa inserts their own local internal TLD in the root zone. Now your root zone doesn't validate anymore. What do you do now?

So what we do need to see is – well, the protocol, when it was designed from a security viewpoint was right. DNSSEC does not work. We will not give you an answer that's secure. But from a user, that's not very helpful, because I want to go to Facebook.

We need to get some way to present the failure of DNSSEC to the user with options of going further, maybe do some probing like DNSSEC [inaudible]. It does a lot of different tests and programs to find out if it can get DNSSEC anyway, present some options to users, say, "Okay, just like with SSL certificates, I'm fine. I'll go [further insecure] or try and use Google DNS," or turn off your laptop and go somewhere else. Next slide.

So for future work, this is all very blue sky. I'm just giving a quick overview. Some of you may have heard me talk about this before. We hope to finally get someone working on this. Better support and application for DNSSEC.

We've been looking at getdns-api which is a very nice first step to make it easier to actually get your DNSSEC information useful in your application, but we actually want to go a step further. Actually, we want to go a few steps back.

If you look at what does the application develop or want, he does not want to do DNS. He does not want to think about setting up sockets. He especially doesn't want to think about [inaudible] for checking signatures. He wants to write his application.

So wouldn't it be great if the operating system would do everything for him? So basically, what your application would do is go to the operating system and say, "Give me a socket to facebook.com and I want it encrypted."

Then the operating system goes out, does the DNS, does DNSSEC, sets up TLS, checks the certificate, goes back to DNS, does DNSSEC, does DANE, and then comes back to the [inaudible] and say, "Here's your connection. Go ahead." Of course you also will need some way of sending the information back to the application to prove that it is secure, or even if something breaks along the way, what is wrong and let the application prevent that to the user. But that would be a really great way for application developers to not actually have to think about all this hard stuff, especially cryptography. They've been getting it

wrong so many times in so many new, interesting ways. So let's do it for them.

There was one more thing I wanted to say about this slide. Ah, yes! For [inaudible] specifically that would be probably interesting to put in the [inaudible] which is a part of the [inaudible] project which is FreeBDS's way of sandboxing. So for applications that don't have specific capabilities, they can go up to the [inaudible] and say, "Could you do this for me?" which also means that the cryptographic functions will be sandboxed and hopefully have one more [inaudible] less in your application. And that's my last slide.

RUSS MUNDY: Thank you, Erwin. That was very interesting. I know FreeBSD has, for many years, been a bit of a forerunner, groundbreaker, in new ways of doing things. I think you continue to do so and that's of great benefit to the overall community.

Next we have a virtual Paul Wouters, also known as Wes Hardaker. Wes is pretty familiar with what the work is that Paul and his folks over at the Fedora group are doing, so that's why we asked him to sit in for Paul since Paul was unable to physically be here himself. Wes, go ahead.

WES HARDAKER: All right. Thanks very much. I'm Wes Hardaker and I'm speaking for Paul Wouters who I talk to on a fairly regular basis. I manage a good percentage of the Fedora packages for DNS and DNSSEC. He does the vast majority of the work and I know a fair amount about what he does. I don't know everything, so I'll answer your questions when I can. I was

hoping he'd be online, but he's not at the moment. Hopefully maybe he'll chime in at the end.

Red Hat has a couple of different products that they really support. Fedora is the lead of the pack, cutting-edge release for Linux. It's sort of considered that by most of the Linux world. They're really trying to push ahead most of the time.

And then of course they have their Red Hat enterprise stuff, as well as they do support CentOS, which is the free version of Red Hat Enterprise without support.

There's really two different sets of package trees that you can hit depending on which version on the operating system you're using. Some things are in one set and some things are another. But in the Fedora and EPEL package trees, there are lots of DNSSEC servers. They're all the ones you'd sort of expect including BIND, unbound, NSD, knot, and pdns.

There's a whole bunch of signers BIND has of course, root signing stuff. And OpenDNSSEC is there and our DNSSEC tools package is in there, which includes our zone signing application.

And then there's DNSSEC utilities which do other related things including – there's a bunch of libraries as well as hash-slinger, which helps you create DS records, and we have one that does that, too. There's some open PGP stuff that tries to tie PGP, which we'll talk about later. Lots of other scripts to help you generate records as needed, like SSH fingerprints and things like that for publishing in the DNS.

There's DNSSEC desktop integration. dnssec-trigger will help you move around to wireless spot to wireless spot, which we talked about earlier.

And then there is some VPN support via unbound and resolv.conf configuration. More on that is – that's the next phase that I think the world at large is looking at is how to tie DNSSEC and VPNs together better in the future. Next.

So in Red Hat Enterprise Linux, there's supported software in their core or collection. So this is sort of the stuff that Red Hat themselves does. And then there's unsupported software, which is still available in the EPEL repository.

And to give you a clue, everything I publish is not officially supported because it's contributed by me and I'm not a Red Hat employee. So everything I do is in the EPEL, whereas a lot of stuff that Paul publishes is directly in their core packages.

So there's a split, but everything is available in one place or the other, depending on whether you're willing to add the extra repository to your system or not.

The servers – only BIND and unbound are in the official Red Hat Enterprise package tree. The other ones are in the other tree, for example. I think that's good. Next.

They're working on cloud support for DNS and DNSSEC. They realize that's a challenge problem. Personally, I don't have a huge amount of knowledge on the subject, but they are working on running hundreds of containers and VMs. They're trying to optimize it. So they don't to run a

validator in each one of those containers. They're trying to look at ways to share a validator across a whole bunch of various containers.

They're working on putting DNSSEC detection directly inside of glibc for at least dealing with the AD bit. They're not necessarily validating within glibc. They're looking at systemd and ways that they can integrate DNS more directly into that as well.

So various types of things. There's a draft I believe that Paul might be one of the authors on, which is the ietf-dnsop-edns-query-chain draft and you should go look at that if you want more information.

And again, they're working on IPsec and DNSSEC support as well.

And then, eventually, they want to move – dnssec-trigger, if you know much about it, is sort of a wrapper around your name server management. And as you connect and disconnect from various things, it will sort of detect using a call from network managers saying, "Hey, I've moved. You might want to go do something to see if you're still in a DNSSEC-capable place." And depending on what the answer is, it'll change things.

They'd like to move that more so it's directly maintained in network manager as opposed to an external third-party plugin, which would make it easier to distribute to everybody, I think. Next.

RUSS MUNDY: Paul had a quick comment. He said most customers will only run systemd plus the application.

WES HARDAKER: Oh, he's in the – excellent.

RUSS MUNDY: He did make it online.

WES HARDAKER: Good. Good timing, Paul. Next. So he can correct me if I ever say anything wrong, which is likely, by the way.

DNSSEC and crypto. There's various cryptographic libraries which are really the only ones that are provided directly by the Red Hat Enterprise package tree that I mentioned before. Again, a lot of other ones are available in the external package tree.

OpenDNSSEC integration is possible due to softhsm. I don't think that there's direct support for hardhsm directly in the main package tree yet, but Paul can correct me if I'm wrong there.

And then for ECC, only "Suite B" curves are allowed in the main tree. There's no GOST and no some of the others. The newer developments in the crypto libraries are not there yet in the primary support. Next, please.

UNIDENTIFIED FEMALE: That's it.

WES HARDAKER: That's it, okay. Again, we can channel Paul, because he is online. He can answer better than most. Thanks for your attention.

RUSS MUNDY: Okay, thank you. There was one quick correction Paul just put in there: "It's Brainpool not Brainstorm curves."

Okay. We have made it to the end of our presentations, so I'll open the floor for questions for the OS representatives here on the panel. Dan, I see you're first in line.

DAN YORK: Quick question. Thank you to all three of you for doing that. Quick question for Kumar from Microsoft. Do you have any thoughts around enabling the DNSSEC validation by default or will it always just be an option that's…?

KUMAR ASHUTOSH: DNSSEC validation by default in the client or on the resolvers?

DAN YORK: On the resolves.

KUMAR ASHUTOSH: Resolvers it is [enabled]. It's just there. You add the [inaudible]. It is enabled.

DAN YORK: Okay.

KUMAR ASHUTOSH: So if you want to say that the root [inaudible] should by, by default, there or not – so we provide a way to…you just write a script. Just type a command. It will just import the DNS root to [inaudible] from the URL [inaudible] published. That's the only step.

DAN YORK: Okay. I was more just curious, sort of to what Erwin was talking about with their plans to move to where they'll have unbound on by default in their 11.0 version type of thing. I think Fedora – does Fedora already do that? I'm not sure. At some point they were looking at that, too.

I'm just curious if Windows server would start to enable that by – if there's been any thought about just making it so that the trust anchor's already there in some way or something so that it would be enabled by default. Just more curiosity than anything.

KUMAR ASHUTOSH: I think I can't answer that question right now, frankly. I think the current idea is the flexibility lies with the operator who deploys that resolver. And it's a very simple step. It's a no-brainer, actually.

DAN YORK: Sure. But it just gets back to that bootstrapping question, because it is a step. So somebody's got to go and download that trust anchor and insert it. And from a security point of view, I can understand, too, the argument that that's perhaps more secure because they know they're pulling down that, they're putting it in there. They're doing that step to

do it. But it is just one more step. I was just more curious if there had been thought given to that. That's all. Thanks.

BILL DRAKE: I see Mehmet's hand up. Go ahead if you want to add or whatever.

MEHMET AKCIN: I just want to actually, not answer because Kumar provided the answer. But these are really amazing suggestions, and if you have these kinds of suggestions so that we can take as a homework and try to deliver as much as possible, feel free to send an e-mail to dns@microsoft.com. That goes to me, Kumar, several people. And this is not a commitment to deliver that, but this helps us to justify. As well as there is an online mailing list just like BIND users. There's a VIM DNS users. Join and feel free to raise these kinds of suggestions there. We are happy to help.

KUMAR ASHUTOSH: I'm sorry to take your time for this, but I think you can just catch me offline to discuss anything that you want. Okay.

RUSS MUNDY: We also have a question in the chat room.

[JULIE]: So this is from [Olaf Rg]: "When can we expect secure connect socket primitive from all FreeBSD and Microsoft? Really like that as connection primitive."

RUSS MUNDY: I see puzzled looks here. That is indeed what [Olaf] asked.

UNIDENTIFIED MALE: Is that related to the last slide I presented?

UNIDENTIFIED MALE: It says “secure connections socket primitive from all BSD and MS.” I don’t know how it is connected to DNS server at large. Do you have any answers?

UNIDENTIFIED MALE: I’m not sure I understand he question.

RUSS MUNDY: So [Olaf] maybe you can put a little bit more in the chat room and we’ll go on to another question and come back to yours in a moment. Geoff, go ahead.

GEOFF HUSTON: Geoff Huston, APNIC. Right now the relationship between the end client and the first hot resolver is basically in [inaudible] resolv.conf. If you list them, you get them via DHCP or something. You send a query. You get back an answer. It may or may not have the AD bit set. You have no idea if the AD bit should have been set. The whole thing seems a bit weird.

Do any of you believe that you’re heading down a direction where the actual end user validates the results they’re getting themselves, so that

they too take the [rrsig] data and they directly do validation. So instead of just one query against their local server, they might do a few to demonstrate to themselves they're happy.

Do you see that as a direction for your products? I don't know whom to ask. I'm a FreeBSD user, so I'm really looking at FreeBSD, but it's a general question.

UNIDENTIFIED MALE: I'll let him answer first then.

[ERWIN LANSING]: So you're thinking of not doing the whole caching validation like unbound, but do something less than that but still…?

UNIDENTIFIED MALE: Don't forget, the resolver has all the answers cached. The entire validation path is actually sitting in the first top resolver. What's going on then is the client asks a bit more from its local resolver, because it's not just relying on the AD bit. As I said, you never know if it should be there or not. It's kind of a hit and miss.

And the other way of doing this is to say to the end client, "Don't trust it." Just keep asking the resolver and drain its cache, and validate yourself and be happy. It's kind of a "Do you see this heading into your client software systems?"

[ERWIN LANSING]: Well, the [inaudible] to do everything locally, no matter what your upstream resolver does. If it does validation, good. We'll do it again, because we don't trust your [inaudible] anyway, unless we set up a [VPN too] and do all that overhead. So we just import unbound and do everything ourselves, no matter what our [parent] does. It will be a lot more [queries].

[WES HARDAKER]: So I had certainly given this a lot of thought and I'm speaking both what I believe the Red Hat approach is likely to go, as well as based on lots of conversations with lots of people, because at Parsons, we have a validating library that we use extensively in all the applications that we write, and we do it in the application.

Me personally, I like to do, especially security-related decisions in the application. So anytime I'm doing cryptographic bootstrapping with DNSSEC, like with SSH or with DANE or things like that, I want to do it directly in the application.

That is not a 100% wide taken view, however, and many people believe – the common middle-ground seems to be that running a locally-validating resolver on local host is sufficiently secure.

For example, I'll speak later about DANE and SMPT and [post-fix] did inside of that, and that's what they're expecting is the system administrator that's running an application like [post-fix] where they are doing DANE must do so using a local caching, local-validating resolver and they'd better be doing the right thing, because there's no way to check for it.

RUSS MUNDY: Okay. I think we have clarification from [Olaf's] question. Julie, did you want to read that again, please, with the longer version?

[JULIE]: Also, before that, there were some comments from Paul that may be relevant to the slides, so if it's okay. Paul said, "Ideally, the hot spot/local DNS is only used to gain some secure resolver from elsewhere on the net."

And then he says, "[Think] throw away the container to log on the network. Then using local resolver only as forwarder and doing its own validation."

And then [Olaf] clarifies, "Applications want a simple, powerful primitive to connect to servers. [Thus] I want the OS/system libraries to take on the work of dealing with DNS answers and SSL connection setup."

And it wasn't still clear to me which slide he was referencing.

RUSS MUNDY: [inaudible] FreeBSD.

[JULIE]: Okay. He says thanks.

[ERWIN LANSING]: So if it's this slide, it's blue sky thoughts on a paper napkin. We really [hope] someone to start working on this soon, but it will be some time off. We don't have any code or design yet in any detail.

RUSS MUNDY: But lots of us like the idea. And we have another comment from Paul.

[JULIE]: Paul says, "Added to Wes and above, the problem is large local networks with dozens of local domains not advertised via DHCP."

RUSS MUNDY: Okay. Do we have more questions in the room here? Any behind me? Okay. Well, it looks like everybody is anxious to get to lunch. We're about six minutes early, so we are actually ahead of schedule. This is excellent.

Dan, before we go to yours, I want to say how thrilled I am that we have three major operating system vendors all actively engaged in doing DNSSEC and incorporating DNSSEC. That is a huge credit to the organizations. I'm really glad you're here.

[applause]

Dan?

DAN YORK: Yeah. Sorry I didn't catch you before in your last round. I did have one question for Erwin, since we have a couple of minutes. Could you just

explain what was the University of Ottawa doing with hacking the root zone or something?

ERWIN LANSING: Yes. They insert their own TLD.

UNIDENTIFIED MALE: What is it?

ERWIN LANSING: I don't remember. But there was an extra TLD in the root zone, so the root zone doesn't validate anymore. There's ways to do this unbound or as a local zone, etc., but [inaudible] directly out of their – made their own root ZONE.

DAN YORK: Okay, thank you. I know it can be done, but I just had not actually heard of it being done. Thank you for that.

ERWIN LANSING: What you'll see is interesting things happening to DNS lots of places when [they go to these], especially the free hot spots during [NX] redirection, inserting java script ads into your things. Of course they also break DNS, right?

RUSS MUNDY: Okay. If we could thank our panelists one more time.

[applause]

Julie, over to you.

[JULIE]: So lunch should be outside. I saw the preparation for it. We were telling them not to put the food out until basically the last minute, so that it did not lure people from elsewhere in the facility, even though we are supposed to have the tickets.

They're not going to be checking your tickets as you walk out, because they know, of course, that you're in the room where you're supposed to be. They are going to be checking tickets of people walking in. And there's also a room over there that's being used.

There are a very limited number of tables that you can stand at out there, so some of you may want to bring your food in here, which is perfectly fine. Apologies for that, but that's just the way it goes sometimes. Anyway, enjoy, and thank you.

RUSS MUNDY: Also, if you are going to wander down the hallway, if you need to go to the restrooms or anything like that or you want to go out there, please do take a ticket so you can come back and get food.

[break]

DAN YORK: Good afternoon. How are we doing? All right, we're quiet. Come on folks. This is the DNSSEC workshop. We're in the afternoon session. First

I'd like to have everyone, if you could, give a round of applause for our sponsors for that wonderful lunch that we had right here.

[applause]

It is those sponsors that do allow us to give all of us food. Again, it's Afilias, CIRA, Dyn, Microsoft, .SE and SIDN. We appreciate their support. We hope they'll continue to support that. I'll be talking to them on that one.

This afternoon we've got a panel that is going on here. Roy's here, and Roy, you can sit right next – there, that's good. It's just Wes's laptop. He's moved over here. We kicked him over into the middle area.

We've got Joe Abley from Dyn, Roy Arends from Nominet who is over here, Wes Hardaker who's in the middle. We've got Kaveh from RIPE on the end, and we've got Jacob from CIRA who's in the middle. So we've got a range of panelists to talk about this issue about DNS and DNSSEC monitoring from a variety of perspectives.

Each panelist has about ten minutes. We've decided to do that. I am going to have questions for the panelists in their time. It may not run exactly ten minutes apiece, but we do have an hour total for this session.

Panelists, Mr. Ebersman over there is running our nice little clock with our countdown timer, although it doesn't do anything maddening when it expires.

We're going to begin with Joe.

JOE ABLEY: Oh, there's a variety of clocks. This is very good.

UNIDENTIFIED MALE: You are being watched!

UNIDENTIFIED MALE: We are going to make sure that it runs on time here!

JOE ABLEY: I apologize in advance to anybody who's expecting actual facts or knowledge to be imparted during this particular part of the presentation. I've decided to try and go for the least-informed presentation for the entire meeting here in LA. I'm fairly confident that I'm going to win the prize.

Having been the last person to supply slides, as Julie notes, and the last person to choose a topic to talk on within the frame of this panel, it turns out all the good subjects are already gone. So I was left basically trying to find something to talk about, the most recent irritation that had happened in my conversation with my excellent colleagues and this is what came up. So apologies in advance. Next slide.

I'm talking about SLAs, and particularly SLAs from the perspective of somebody who is most decidedly not a lawyer, but is sometimes called upon to help one side or the other with a conflict that results from an SLA.

You can have your own opinion about the value of an SLA. I think sometimes we think it's unnecessary, it's an extra piece of contract that

will never come to anything. But chances are, if the text is there, at some point somebody will be encouraged to use it to break a contract, to change something, to negotiate better terms, to get a refund on some money they paid in the commercial realm.

Perhaps this is true of all elements where SLAs are potential vehicles for buying service from someone, but I think in the DNS in particular, we have a particular weakness for understanding how to write SLA text, and the industry being somewhat newer than some of the other Internet services you can buy. So let's skip ahead to the next slide.

This is the kind of thing that drives me insane. We end up with language in contracts that say things like there is a query latency requirement that must be less than some figure. I put 200 milliseconds here, which is a number I just made up. And someone decides, "Oh, query latencies are higher than 200 milliseconds, so we need a refund or we need some action to be taken."

Then it goes down to people who say, on the other side, the question then becomes, "Are we really over 200 milliseconds? Is something broken?" And then you say, "What do you mean by 200 milliseconds?" And then you just get blank stares. 200 milliseconds, what does that mean? Are we in compliance? Are we not in compliance? What did we even mean by 200 milliseconds to start with?

For this particular example on the next slide, examples of where you might see reliable latencies for queries greater than 200 milliseconds. There's an infinite number of scenarios where your query latency is going to be high and you don't have to live on a Pacific Island. Having a teenager in a bedroom with a laptop is enough most of the time to

make sure that everything you can measure is more than 200 milliseconds.

And this doesn't make any sense. There's no point putting this language into a contract, because it means nothing and it's just a recipe for arguments. So next slide, which I think I probably just covered. There we go, yeah.

This last one, I would like this to be a quote from someone, but I just made it up last night, so you can feel free to quote me. I think the bear trap of ambiguity has a t-shirt written all over it there.

This is the basic problem. People specify what they want in terms of service, but they don't understand how to write it down accurately. Next slide.

The title of this presentation was deliberately inflammatory, because we don't expect lawyers – the problem here is ambiguous language and the job of a lawyer is to avoid ambiguous language or to exploit it if it already exists and you're trying to get an outcome of a particular kind.

The problem here is not with the lawyers who are drafting this text. It's just that they're asking the wrong questions and they're certainly getting the wrong answers. Someone at some point saying, "We need an SLA," and someone else is saying, "Well, what kind of SLA can we give them?" And then they say, "Well, query latency is something they mentioned, 200 milliseconds sounds like a good figure. Can we meet that?" "Yes, we can meet 200 milliseconds. That's easy." Okay, we'll put that text in.

The fault here is not the lawyer. The fault here is that the wrong question is being asked and the wrong answer is being given and the context are completely different. So the lawyer is looking at text that can be defensible in an argument, that can push things in a certain way. The technician is not even vaguely thinking about that. The technician has got a particular mindset and a particular set of graphs and graphite that show a certain thing and they know that 200 milliseconds sounds easy, because graphite [inaudible] around 20 milliseconds, so that seems fine.

But nowhere in this is any appreciation for the fact that there's an end user service somewhere that has to reach a certain acceptable level and no one really knows what that level is. Next slide.

So as inputs into this, there are some things we have to appreciate. The Internet is littered with other people's networks, which depending on where you work may be better than your networks or they may be worse than your networks. But it's a seething chaotic mess. It's a lava pit filled with – I don't know – prehistoric lava-dwelling fish biting each other. It's not something you can easily characterize. This is not a piece of fiber with standard transmission characteristics that you understand, where you understand what a fault means. This is a ridiculous network that should never work in the first place, and the fact it does is, frankly, a miracle. And this is something we're trying to make measurements over.

The other point down here, which is my absolute favorite definition of the Internet ever – most accurate, and yet useless simultaneously definition that you could possibly get – is my Internet is not the same as

your Internet. Even on the same wireless network, I'm probably getting a different experience from you, but certainly from different parts in the typology. There are different parts of the Internet that are missing.

There are black holes that exist for me that don't exist for you and some of this is just a natural consequence of the kind of way that routing information floods in chaotic spurts across the network. Some of it's to do with policy filters. There's firewalls that are here that are different from somewhere else. There's NX domain redirection. There's domains that are being [written] on the fly. There's captive portals. Everybody's Internet looks different from everybody else's Internet, so my 200 milliseconds is not the same as your 200 milliseconds. And you can't write SLAs without defining exactly what you mean. So, next slide.

So here are some ideas. I have never seen a good SLA, so I don't really know how to tell you how to write one, but here are some ideas that I came up with after considerable thought last night and five minutes while I wrote the slide deck.

Good SLA language is tied to tests that are reproducible and precisely described. What you want here if your goal is to improve service or to reach a common understanding with the supplier, you want to be able to both have a visibility of the same numbers and both agree over the nature of a problem at the start of a conversation, because if you don't have that, your conversation is not going to end well.

You need to have measurements that understand that the Internet changes and it's other people's networks and you don't have control over it and you have to be reasonable and understand the occasional

functionality of the Internet is something to be pleased about and not the [inaudible] to be angry about.

And the last thing I think is the key point. You might imagine that query latency is the be-all and end-all for a user's experience and that's a reasonable thing to measure. You might imagine that there are many other things, but you're not looking at what the actual user is trying to do. You don't have an understanding of the end user. You don't understand what acceptable means. You're guessing at metrics that are apparently easy to measure, if they're well-specified, but they're arbitrary and no one has really done the thinking to understand what is it that we actually care about.

Do we care about the system as a whole across a set of authority servers? Do we care about resolution as viewed through a cache which might lead you towards TTLs and other things? There's many, many variables here beyond these very simple course measurements that just say, "This is acceptable and that's not."

So to be honest, an SLA that specifies a threshold of 200 milliseconds, it's a better contract if you just cross that out. It's not buying you anything. It's just making your future life difficult. Next slide.

So some ideas that have come up recently. How can you get around this for people who have been bitten before by problems with previous suppliers and have management that insists that the next contract will be better and we won't be stuck like this again?

First of all, as I said, understanding of what acceptable service looks like is fundamental, and understanding of what kind of failures and their impacts is fundamental

Third-party measurement systems are an easy way to not define a measurement that you want to do. For instance, saying round trip latency to these particular servers using catch-point service or using a particular range of tests running on Atlas probes or something else like that, that's something that you can both do and you can both agree on what the results are, and then you can have a conversation about whether or not the service is adequate or not. I used the phrase "throat punch." I was quite pleased with that last night, and then I went to bed.

So really, what I'm hoping for here – and we skip to the next slide – is really to have a conversation and to get people thinking about these things and not just randomly throw arbitrary measurements into contracts. It would help me enormously as a provider of commercial DNS service if I could get some ideas from people here as to new things we could measure and new things we could judge service quality on.

Any feedback here, if this resonates with anybody at all, I would love to hear more about what things you think work and what things don't work and what problems you've had with SLAs in the past, because as one commercial DNS provider, we can at least do our bit to try and make these SLAs more reasonable in the future and stop making engineers have to quietly weep in corners whenever the lawyer walks into the room. Thanks.

DAN YORK: Thank you, Joe. The reason we bring you on these panels sometimes, Joe, is at least you provide the entertainment value. There's a level of poetry in here that sets a good bar here. Anyone have questions for Joe about steaming pools of other people's networks or things? I see Martin coming here.

[MARTIN]: I do think it's about recognition of the first use of the phrase "throat punch" in a presentation at ICANN is worth something.

JOE ABLEY: I think you have to go to the records and transcripts to prove that. Your data pour in this whole presentation, so that's your homework. Entertainment, yes, but let's get here a serious point. You have two different points you brought up: my Internet is not your Internet, which is actually an amazingly ill-understood issue but has been around for a long time. I'll take your arbitrary 200 milliseconds. You measure 200 milliseconds and say my SLA works; I measure 200 milliseconds and it's the other way around. I say the SLA is in breach. We have no resolution for this.

You're pitching a good point, but I think we have to move away from the classic SLA thinking that we've dealt with, at least in Internet transport, which is where I focused for a long time. I think you've got to really throw this on its head and come up with a whole new way of measuring it. Otherwise, we'll be arguing the wrong thing for a long time. Food for thought.

DAN YORK: Okay. Warren?

WARREN KUMARI: It seems like a lot of this can also be avoided by simply taking the network portion out of it so that you will always return responses to queries within X of receiving them. Which is fundamentally useless trying to claim against it, but at least it helps you understand what's going on.

JOE ABLEY: It's true. Like everybody else here who has commercial DNS servers, these kinds of metrics are things you track to try to promote, internally, the health of your service. No matter what contracts you have today, if your service goes downhill and doesn't stop, at some point the customer's going to go away. That's an example of a measurement that happens inside the network which is not reproduced by outside the network. As an SLA gauge, it's also useless.

DAN YORK: With that, I would like to move on to our next panelist. We will come back at the end for some more questions around this. Hopefully to offer some definitions of how we can do this and what we can measure that would fit into SLAs, I've got first up is Roy Arends.

ROY ARENDS: Thank you. Let's go to the second slide immediately. I was asked to think about monitoring DNS and DNSSEC. This is the same time where we at Nominet do a lot of monitoring, do a lot of data analysis, etc. T

his brought me back to a couple of years ago when we had incident at Nominet. We had DSC graphs and we could see that something was happening. We just didn't know what.

Immediately we thought about what can we do to actually see this? So you turn on [inaudible] and you wait for a while and hope the incident happens again. Luckily, it did, so we could analyze what was essentially the problem. We realized that what we were missing is the ability to analyze all that data. I've already given a presentation about a tool that we've been using earlier this week, so I don't want to go into that.

I just want to go into what you can actually see when you look for things. Just like any other IT department, we have an IT department that has things like Nagios. Nagios will do some threshold measuring. If stuff goes over 80%, you get a phone call or an SMS message, etc. You can measure things like disk and CPU and memory, etc. That's all very interesting from an administrative point of view, but you still don't understand what's happening in the real world actually with the traffic.

When we started at looking at what was happening, instead of just knowing that something was happening, we saw a few interesting things. I've used this time to basically tell you a few stories of the stuff that we've seen. Can we go to the next slide? Can we skip this slide and the next slide? Thank you. These are the case studies that I tried to fit in these ten minutes. Next slide. Thank you.

I apologize to Warren, I'm calling him out again. A couple of years ago, as you know, Google started deploying 8.8.8.8. When they did so – go ahead, we'll wait.

When we looked at traffic on average we saw that the amount of [surfels] go up. When we just built this system we had a trigger in it. When we see a SERVFAIL, we want to raise an alarm because SERVFAIL is bad. [Surfel] is basically a temporary failure and you know that stuff is going wrong. This is on our authoritative servers. These phones were going off all the time. When we looked at where the traffic was, this is basically because of Google sending us a very, very long string.

What Google is doing is actually very smart. They're prepending a query name with a nuance. They basically wait until they see something like an NX domain and then they know they are at the end of their resolution path. They do this to increase the amount of entropy in a session. Basically a query, a response.

However, you still need to see the difference between a wild card response or an endless delegation and an NX domain. If there's a wild card, you'd never see an NX domain. Another small thing, they prepended that whole string yet again with a nuance. Of course, you've still got a no error and prepended with a nuance. So this was a loop inside of this query stream. Slowly, Google were increasing their string and it became over 255 bytes. A queue name over 255 bytes is actually not a legal queue name by convention. Not by any limits of the system, but by convention.

As you may recall, a domain name is a set of labels. These labels have a specific length, but the amount of labels is not a specific number. Google was just happy creating longer than 255 bytes and that results in another bug in BIND would return SERVFAIL. BIND should have responded with form error as in "I can't compute," but it was SERVFAIL.

Google thinks, “Hey, a SERVFAIL! A temporary failure! Let’s do the exact same thing on the next server.” Which would eventually return into a SERVFAIL, the exact same thing on the next server, etc.

That’s why we see these large peaks. You can’t actually see the – this little circle, if you look at this, it’s at six [inaudible], but that’s because I’m filtering on a specific IP address from Google. At the same time, we informed Google and [ISE] and it was fixed pretty quickly. It’s not really something that is really, really bad, but it helps Google [live] at faster resolution, it helps [ISE] to have a little bit more stable products. That was the Google backstory. I’ve got another one. Next slide please.

OpenDNS. OpenDNS, as you know, is an open resolver. They have a lot of interesting things going on. One of the things they do is they limit their shutter time. The shutter time is nothing else than the amount of time you wait for a response to come in. If you know a little bit about resolvers or clients behind resolvers, if they don’t get their drugs they become very, very, very aggressive. They’re like junkies.

Here we have a case where the distance – let’s use the distance in milliseconds (I apologize for the lack of purity here) – the latency was 160 milliseconds between, I think a machine in Singapore and NS5 or NS4. One of our name servers. They were actually waiting for 300 milliseconds for a response. But if the latency was under 60 milliseconds on average, the return trip is 320, which is 20 milliseconds after they close the window. After 300 milliseconds, they ask again. And again, and again.

You can see this is a pretty massive spike. This is a couple of, maybe a couple of hundred standard deviations over what we normally see from that IP address.

OpenDNS use for their resolver site, similar as Google has, many different IP addresses. Even though on the front side you see 8.8.8.8., the backend has many different IP addresses. This was just one IP address that created that amount of traffic. That's yet another thing we saw with this tool. Can we go to the next slide?

This is a long time ago. This is about three years ago. We know the behavior of a system. We know that we should never respond with anything related to dynamic update because we do not allow dynamic update. If we [normally] get a dynamic update, we send back refused, as many others do. We should never see anything else than refused. You can actually measure for that. You can look for that. You can see [inaudible]. If you see something like NXRRSET, you should be very scared because you don't allow dynamic update.

When we saw that, we tried to recreate that. This package that you see on the screen was actually fairly benign. But since I'm not really a good developer and when I try to recreate the packet, I crashed my local machine. It turns out I was using an old version of BIND, so I upgraded that and crashed that too. This is already a long time ago. Your system is safe now. This was in 2011, an old version of BIND.

This thing's actually analyzing this and with some pretty straightforward visualization you can find these things. This stuff is important. For instance, at the time we had the bulk of our name servers were BIND and a few other name servers were NSD. We calculated that if we lose

half of our servers, the other half would basically see the traffic of that first half. This would be pretty significant. It is actually pretty bad. As you know, BIND is very prolific. A lot of people run BIND. It is a problem. Next slide, please. I will be quick.

This is a botnet that we found using this. A lot of the top level domains see this. You see a lot of MX queries. If you look further, they have the RD bits set. If you look further, the resolve off the NXDomains because people have a lot of e-mail addresses that are false on websites. It's to dilute spam sets. Using this tool we could actual sieve out a very, very specific fingerprint. Not the entire fingerprint is visible on the screen, but come and talk to me if you're interested in this.

Therefore, we have now instead of IP addresses that talk to us directly in order to send spam, which you can then hand out to other companies in order to block that spam. Next slide, please.

The index case. This is a very complex story. This slide is available, I'm going to skip that, I have 20 seconds left. Next slide. This is yet another example of badly deployed name servers. This is an IP address that our tool immediately found not being random. You can actually look for that, but now you see that it is not random. We see this all the time. Half of the stuff that we see has some problems with it. Half of the IP addresses that we see querying us have a problem i.e. no random identifiers. Yes, not just port numbers, not random identifiers. Also random port numbers, not random port numbers. We still have name servers [inaudible] on port 553. I have no idea what it is. It is DNS traffic, but we don't respond to that. Next slide, please.

Take-up of IPv6 and DNSSEC. You can look over a couple of years, see the uptake of IPv6, uptake of DNSSEC, etc. Next slide please. I'm greatly over time.

This is why analysis is important. You can do monitoring and you will see incidents. The important part is that you actually investigate those incidents. That's why analysis is important. You can see that often when someone screams DDoS, it's actually just a misconfiguration.

At Nominet we leave monitoring the health of the system to very good tools like Nagios, but the analysis, we have basically built our own thing. Next slide please. Great. I'm done. Thank you.

DAN YORK: Roy, you've actually still got – oh, that's over. The counter went. Okay. Well, in light of the time, because I want to make sure everybody gets in here, let's park the questions. If you've got any for Roy hold onto those and we'll see if we can come back with that.

I think we're now going to go to Wes. Do you want to do Wes and the demo? That's the plan. For the people who are remote, we're going to be switching to a demo for the next few minutes while Wes does this. It's not visible to you out there.

WES HARDAKER: I'm going to do the slides and the demo both. I'm going to do the slides and the demo both. Let me just start by saying I'm going to do something stupid and I'm going to start by doing a live demo. Better yet, I'm going to do a live demo of somebody's domain names. Somebody

give me a domain name that you don't mind me sending a few queries to. RIPEe.net.

I'm going to add this really quickly. This is our DNS sent no monitoring package which we're really just launching and publicizing for the first time at this conference. It's going to go off and look for a whole bunch of stuff to monitor. Look at all the stuff it's found! We won't go into it, we'll just hit submit because it's all good. That way I can talk about the slides more. We'll come back to this in a minute. I'm going through a quick setup of a new domain name. We'll come back in a minute and look at it. Back to the slides, which was there.

The fundamental point that I want to get across today is when you're monitoring you DNS service, where are you doing it? If you're monitoring it from your [NOC] and your home location, that's not where your users are. We have a new service that we're offering that will monitor it from what's the perspective that your users are seeing? What's the entire Internet see of your domain? We think that that's a critical thing to think about.

Your customers need to know whether your name servers are up, of course. Your customers need your name servers to be serving the correct data, and consistent data. They have to be equal. But what do your customers see? You probably know what you see. If you measure it from your [NOC], you see it very quickly and you see it without any hops and bubbles along the way. They're somewhere else. Where are they? Is your zone consistent everywhere across the world? Are you able to measure that?

DNS Sentinel, which we're sort of in a beta phase at the moment and we'll be rolling it out to some beta customers, is a distributed DNS monitoring platform. It detects and reports all sorts of stuff, including changes in zone data, DNSSEC related errors and warnings, DNS errors and consistency issues. We can tell you if your name servers are aligned. We track WHOIS changes to see if your WHOIS information is consistent from one moment to the next. That's a common source of attack, which is one of the reasons why we look at it.

We do the standard timing analysis too. We can show you a graph of how long it takes from all sorts of places in the world to your name servers to give you an indication of how your zone is performing according to the user's view. There's many more things we have planned on the to-do list, but we're now at the point where we can at least roll out the minimal things that we've implemented.

Roughly this is the architecture. If you imagine yourself having two DNS servers, in the upper left and upper right hand corners, we have sensors around the world. That's some five example sensors. They're not necessarily the real ones. We're going to be growing our sensor network to upwards of 20 to even 100-plus sensors around the world so that we can even get a good view of [Anycast], for example.

We have maybe one in South America, Asia, United Kingdom, USA, Germany – all over the place to send queries to all of your name servers.

Then when we detect things, we notify you of problems that we see and events. The goal is to let you customize what's important for you and you can decide what goes to e-mail.

Did you know your technical contact has changed in WHOIS, for example? Did you know that your zone is about to expire? Maybe you just get that via web notification because it's seven days out. Did you know that your website address changed in South America? It may be that the address only changed in one location. Do you know that in this country you're getting a different A record or quad-A record for your website?

This is really the question. Your users are looking at your domain from probably an entirely different place. That's a shot I took in London at the last ICANN conference of a whole bunch of houses. They're somewhere up in that bubble and it's a very different location than what you're looking for. Please, I'll be around for the next day or so and Russ Mundy, who's over there, will be around. If you want a longer demo than what I'm about to give you, we'd be happy to give you one.

One of the things that we would like feedback on is we want to do some public good as well. If there's particular aggregation types of things that we can offer back to the public to give a health metric of the Internet, we're thinking about producing some sort of report that can come out of that based on all the data we have available. If there's specific points of interest that you might want to see, please do let us know.

This is our example. You can see that I actually have four domains listed in here. I'll walk through Parsons.com, which is a company I worked for first. There's all sorts of various things that we can report. Under the status page, for example, you can see that we're checking all sorts of stuff. We can tell you that your parent DNSSEC validity, for example, is valid from October 15th to October 22nd and it's seen by all of our

sensors. It's seen in Singapore, Frankfort, London, Palo Alto and Chicago. They're all agreeing. That's a really good thing.

You can see that there's 292 validated records. That's actually being validated directly on the sensors themselves. They can actually validate all the way to the root. There's some indeterminate, unknown things. I won't go into the details. I could explain it, but we don't have the time. It's actually correct. That's the important thing.

A number of other things. We see mostly we're consistent, but when they're not, we can see the name servers are slightly different from one place to another. That's actually a bug in the database because it's been cleared out. Again, this is a beta system. If we go back to – let's go to an interesting one.

[Flux] is our test zone. We resign this once an hour, we change data in it frequently, we change A records in it, we add and remove stuff, it's our test network. We can do some interesting things with it. Because we change the data on a regular basis, we can actually show you how frequently it's changing. This is actually a visualization of how frequent the data is changing within the zone. You can see that we are actually changing it once an hour. I can make the blip smaller so that you can see that's a periodicity to how often the data underneath the hood is changing.

There's this blank spot where our script actually failed. It didn't change for a long time. Let's go back to RIPE.net. This is RIPE. The system's already up and running and gathered information about it and we've already verified that there's 123 validated records. We haven't found anything that's bad so far. This is actually bogus. You can see that we

found zero errors in a bunch of places. The WHOIS information hasn't populated out yet. It will within the next five or ten minutes.

We can actually begin to look at all the zone data itself that we've collected and you can see that from this sensor in Singapore we're getting all of this data and we can actually go in and look at it. We can see that the validation result is good, where the answer actually came from and all sorts of stuff. The whole point of this system is to let you get a good notion in your head of how your zone is behaving. Not just from where you are, but from around the world.

Who was at the demonstration that I did on Sunday at the DNS-OARC? Okay. You know that it didn't go off quite as cleanly that day as it did today. Somebody said, "Why aren't you monitoring the DNS?" Because I was typing in a DNS name that did not work into the browser. Our system did catch it. The thing is, I hadn't looked at our system, so I didn't notice.

But down here you can see that the address record for beta changed from one address to two. Well, there shouldn't be two addresses. This is where my mistake was. I didn't add the address to the alpha record, like I was trying to demonstrate that day; I added it to beta. So suddenly beta got two addresses.

So our system actually did catch the exact same problem that I hit that day, and then later on when I went and fixed it, it caught the inverse where I dropped one address back. That's the good news. Does anybody have any questions? There's a lot more that I could show and demonstrate here, but there really isn't time to go in, because I have 57 seconds left. But I'd be happy to take some questions about it.

Again, if you're interested in contacting us later, please do find Russ or I later today. Unfortunately, I have to leave on a plane early tomorrow morning, but Russ will be around if you're interested. Again, we're interested in things that you think we could do from the public good. Obviously we can't release client data for those that want to sign up to our service, but we can aggregate and we can produce some data that we think would be a good health metric of the state of the DNS [for] the world at large.

DAN YORK: And of course you were [asking for] questions and you ate up your time there. We've got one over here.

WES HARDAKER: And 21 seconds to go.

UNIDENTIFIED MALE: [inaudible] from New Zealand Registry Services. Are you familiar with [inaudible]?

WES HARDAKER: We know of it. We haven't studied it in great depth.

UNIDENTIFIED MALE: We use it. It reminds me a lot of why you're doing it. You have some [inaudible] features that look very lovely. Are you going to make the locations of your [inaudible] available, publicly available?

WES HARDAKER: We have thought a lot about that. Our plan at this point is not to hid them. We don't see a need to hide them. That being said, that decision has not been made fully yet, so I can't promise that at this point. But I don't think so.

UNIDENTIFIED MALE: I don't mean address, but I mean, for example, physical location and network and [AES] and [inaudible] good spread of [inaudible].

WES HARDAKER: So we'll certainly make it available to our customers, because they need to know is [their Anycast] network actually being appropriately monitored? And unless we can get in a whole bunch of locations to monitor a big [Anycast] cloud, you wouldn't actually know that. We don't want to do [Unicast], because we want to measure the actual [Anycast] performance itself.

UNIDENTIFIED MALE: Exactly that was our challenge.

WES HARDAKER: Yeah, exactly. That's my goal is that we should have that information available to you. Yes.

DAN YORK: Well, thank you, Wes, for this. People will be around and you and Russ will be here. I do appreciate the fact that up on the screen I'm seeing IPv6 addresses. That's awesome.

WES HARDAKER: Absolutely.

DAN YORK: We did not pay for the segue that makes a whole lot of sense here. But our next speaker, Kaveh, is coming at us from the domain that was just monitored in here, RIPE.net. And he's going to talk to us about sensor networks and things that have all sorts of nodes out there. So I'll turn it over to Kaveh.

KAVEH RANJBAR: Good afternoon. I'm going to quickly introduce RIPE Atlas, which is an open public measurement network run by RIPE NCC. For those who might not know, RIPE NCC is not-for-profit membership organization. Its main function is to work as Internet resource registry for Europe, Middle East, and Central Asia but we also do a lot of other projects. We are also a neutral organization. It's important because I will explain about that later. Next slide, please.

This is the probe coverage for RIPE Atlas. Today this morning I checked. We had 6,950 probes up and running. The number at the moment is between 6,800 and 7,000 depending on the connectivity of the users. Next slide, please. Next one, please. Skip this one.

The probes that I'm talking about are the small hardware you see in the middle. There's a TP-link router. We upgraded the firmware. So basically, users get them, connect use before power supply and connect the Internet. It doesn't do any wireless thing or anything. And basically they can forget about it. It does built-in measurements, which includes a lot of DNS measurements as well to all root servers and multiple well-known resolvers and all this stuff.

And also, other users on the network no matter if they [pass] the probe or not, you just need an account which you can create for free. They can use the network to measure different things.

The last one, the [inaudible] unit that you see there, is a bigger kind of probe. It basically has the same functionality, but it's designed for data centers. We know they're more stable because the small probes in many cases are in homes or small offices. Those ones are hosted in data centers, so they have more reliability, better [power], better connectivity. And because of that, we [inaudible] lots more built-in measurements from them. Right now, we have 79 of them installed all around the world. And these 79, first of all, they do [mesh] measurements between each other. They do trace routes and pings to each other every 30 seconds. And we store all this data and the data is available for users to look. Next slide, please.

These are some statistics. I basically mentioned them. Please move on.

So when a measurement is done, no matter what type it is – in this presentation, I will basically focus on DNS measurements. But when a measurement is done, users can download the results. They can [either revisualize] most types of results. But the raw data is available in

[inaudible] XML. We also have Open API. So basically all the functionality that you see on our website, which is atlas.ripe.net by the way, you can see all the functions on the website is built on top of the API. So basically, everything we provide is available through our APIs.

We even have bulk APIs when you have big measurement [inaudible], because some of the measurements have gigabytes of results.

We plus a lot of people – active people in the community – develop different libraries to parse this data, because as I said, sometimes we get huge amounts of results back from system, to parse them to be able to parse them in a language that you prefer. We have published [inaudible] library, but people in our community have published [inaudible], java, .PHP and other [inaudible] for different languages. Next slide, please.

So focusing only on DNS, RIPE Atlas does basically four kinds of measurements. It does [inaudible] certificate checks, DNS, [ICMP], and trace routes. We deliberately do not do any other measurements. We can discuss it offline if you want to know why.

But DNS is one of the measurements we do. And actually, we have a service called DNSMON, which you can access at dnsmon.ripe.net. For DNSMON, we check all name servers in the [root] from the anchors that I mentioned from those 79 anchors all around the world. Every 30 seconds, we send different types of queries with [MSOA] or UDP and TCP and we keep all the results and we visualize them.

So this one is for K-root, and you see on the left you have the instances, and then you have the time. You can go back in time. As I said, we keep all the data. So the data is not aggregated as time passes.

So Atlas measurements, if you have done a measurement three years ago, you can still go back and see the exact results from that measurement, either ping or traceroute or whatever.

Same goes with DNSMON which uses Atlas for measurements, so you can basically zoom in focus, and select any of those single boxes which are a combination of all the measurements for that point in time for a specific servers, and then you can get the raw results. You can see the trace route [inaudible] and then ping, and you can get the actual [inaudible] where we get from UDP or TCP in the DNS query.

So DNSMON is only for 30 ccTLDs, which are mainly operating in our service region and all root servers. But we are adding the same functionality for people who want to set up their own measurement. You can say pick 500 out of those 7,000 probes and measure this domain, mydomain.com. And you will basically see the same format for presentation of data, but it will be basically for you. Next slide, please.

So measuring DNS. As I said, RIPE Atlas measures DNS in [DNS 6]. You can basically, using either probes resolver config or you can define any resolver you want and there's a good reason for that. You can basically send different types of DNS queries and get the raw results back.

Why we let people to change the resolver from the local one is because actually previously we didn't, but after a while we saw that, in many cases, there are people who cached DNS traffic and [played] with it. So

we had basically a national incident in Turkey which we documented also on our website. People were grabbing the traffic and changing the results [on the fly].

So now you can actually, on any of those probes that I showed you all over the world, you can say either use a local resolver or go to another resolver, like go to Google's resolver, but then you can check the results if you see what you expect or if someone in the middle is playing with the results.

We have multiple query types I will show in next slide. On the left, this is where you can set up a DNS measurement. As you can see, you can define – and you have different page for TCP. No, actually, you have a [tick] for TCP – sorry for that.

You can define all kinds of metrics. And on the right, you see part of the choices you have for DNS. So as you see, you have [NDNS] or you have DNS Key, and then multiple other ones.

You send a query from any of the probes or any group of the probes. You can even go by regions or by – nowadays we have tags, so you can say probes who have IPv6 or don't have IPv6, but send queries for [quad A records], we have some of those.

So you can choose any combinations of those 7,000 probes and any number of them you want and define this measurement and say it should run every 30 seconds, for example, for a month. The results will be available.

And we visualize DNS as I showed in DNSMON case, [DNS resource]. But we do not analyze DNSSEC results, at least yet. So what's possible and

what actually has been done is – because we returned raw data, you can do the DNS Key, for example, or any other records you're interested in for DNSSEC measurements and then analyze them later.

Actually, from NLnet Labs, they had an intern I believe – Nicolas – who did exactly the same thing. He was doing it with a controlled name server. So he was also monitoring the queries that were reaching the name server. He released a code public labels presented in the previous ICANN meeting by [Vilam], another colleague from NLnet Lab. The results are already available. The [core] results are available to parse DNSSEC results.

But please feel free. If you want to add anything to this network, this is a public network all for the good of Internet. We also encourage you to join the network. And if you can, you can help us. If you want to receive probes, you can contact us. The contacts are in the last slide. Next slide, please.

So these are the results, for example, from NLnet Labs. This is the last page. The actual presentation was online. You have the link in the presentation. This is their findings. You can set up all kinds of different measurements using DNS and DNSSEC measurements, using RIPE Atlas. The platform is there.

So for us, we basically see RIPE Atlas as a platform. We tried to add features to it, but we really encourage users to use this as a platform. So please feel free to add any kind of analysis tools you want on top of RIPE Atlas network.

This is the contact information. The website is atlas.ripe.net. We have an open mailing list, ripeatlas@ripe.net. And there are multiple articles about Atlas on a publication that we have, labs.ripe.net. Thank you.

DAN YORK: Thank you, Kaveh. And thank you to you at RIPE for running such a great network. I have a probe in my basement and I've done this for a number of measurements. I think it's great. Do you have probes here for people if they're interested?

KAVEH RANJBAR: Actually, we ran out. Yep. I might have some. We run something called Ambassadors.

DAN YORK: What are you talking to me about? That was the classic deer-in-the-headlights look.

KAVEH RANJBAR: I had a couple of them, but I forgot to bring them with me. But you can easily get them. If you write to us, we will ship one to you. But we have this program – actually, I think it's good for this room. We have this program called Ambassadors, because many of us travel a lot and see a lot of the same people with the same mindset.

So if you think you can help us and expanding the network, we can give you 5-10 probes, and then if you go somewhere, you can give them to

the people. And actually, you also earn some credits for giving those probes out.

DAN YORK: Sounds cool.

KAVEH RANJBAR: I already packed, but I forgot to put them in my bag.

DAN YORK: And I would also just mention, too, that what Kaveh mentioned there, that you can do this measurement of DNSSEC about that. Here's what I throw out to the community. I'd love it if somebody would be interested in taking that code and running it on a site somewhere to provide an ongoing trend line of what this looks like.

Because we've got Geoff's great measurements for DNSSEC validation that's running on APNIC from his flash-based advertising the network that he's got, which is cool. It would be neat if we could get a validation trend line coming off the Atlas probe network that was similarly running every day on somebody's site that we could point to to do that kind of validation.

So I throw this out there. The code's out there. Somebody just needs to write that, visualize it, and put it on some site and it would be awesome if somebody would do that. Geoff, what are you going to say?

GEOFF HUSTON: In actual fact, we report from APNIC the flash experiment and the flash servers are actually across the entire globe. Oddly enough, everybody watches YouTube, which is really cool.

DAN YORK: Yes, yes. It would be interesting to see this, too, because there's a delta. Everybody watches YouTube and everybody does things in this that does flash-based stuff, so you're getting a very broad-based thing.

Obviously, the Atlas probe network, it's a little bit oriented toward more of the techie network side of the people who will go and put them in their homes, although I know a lot of people put them in their parents' homes and other places like that, too.

Anyway, I'll throw that out as a challenge for the community. If somebody is interested in that, it would be awesome to do. Thank you, Kaveh, for that and all you're doing. We have one last speaker, Jacob. And if we have any couple minutes left, we will throw out some questions.

UNIDENTIFIED MALE: Can I make a question? Actually, it's a note for the presentation. For those interested in [distributed] monitoring, there is some code part of the Atlas system for getting the results and producing [inaudible] using that. That's really useful for people running DNS servers that don't have the money to pay for.

DAN YORK: It is. We want to bring this monitoring panel to a close with Jacob taking a different spin on what they've done to do monitoring inside of .CA. So this is more of a case study perspective. Take it away, Jacob.

JACOB ZACK: That's right. Again, I'm Jake from CIRA. We took a very literal approach to the request to present here within the DNSSEC group. DNS monitoring at CIRA. Next.

So this is our signing solution. We decided to go with multi-signer. So we sign with both BIND and OpenDNSSEC. Each one of the keys on the signers would indicate a separate HSM that was prepared during our key signing ceremony. So they all contain identical key material. Next.

Coming out of our Oracle cluster, I'm only going to mention basic monitoring stuff once, I promise. RAM CPU, everyone does it. Nothing exciting here. Next.

Once we get into our zone file generators, we are alerting on whether or not we can connect Oracle and get data from Oracle. At the end of our zone generation, we are checking for lines of difference as well as total file change size.

At CIRA, our number one concern would that people far smarter than us had already done DNSSEC and already had a few outages occur, and we basically decided that we would not go under production without protections. We would not go full production without guaranteeing that we wouldn't make the same mistakes and be bitten by the same problems that our predecessors had. So yeah. The too many lines of

difference and too much change in total [inaudible] size, valuable checks to us.

That being said, it [stripped] three or four times in the last year. Each time it was just a registrar doing a whack of changes. But that being said, it's always comfortable to know that we're protected in that way. Next slide.

So, the signers. Lots of stuff here. Obviously we're checking whether or not we can reach the HSMs. We need to ensure that, at all times, the HSMs are in synch. Every transfer of a zone file is atomic. So we're sending the zone file. We're sending an MD5 of the file after the fact just to be able to confirm that the entire file got there to the other end.

So on the validators, this is where a ton of fun stuff happens. On the validators, again, atomic transfer. So we're checking to make sure that the file is there. We are running a valid DNS.

It's important to note I guess that in our signing processes and having taken this method of using two separate signers, we found multiple bugs in BIND and OpenDNSSEC and even valid DNS in how certain things were handled and/or handled differently, and perhaps made the development teams underneath those organizations recheck and reaffirm some of the assertions that they had made in their code.

It's also important to note that KSK rollover at CIRA, while it can be automated, it will not pass our level two validation here within our zone validators without a manual step.

We decided that we will never allow a zone to be published for a [KSK key] if some random bug happens inside [inaudible] cache database. So

we actually keep a file. That's the [DNS key good]. As long as the DNS key set is matching what's in that file, then the zone can be published. Along those lines, we strip out all signatures for the comparison. Next.

So once it gets to the hidden master, we're doing the basic things. I imagine everyone that's publishing zones on a chronological basis is doing. We're making sure that we're actually getting the zone by a scheduled time. Next.

Secondary name servers. Again, same thing. They have a time that they must receive the zone by. Next.

That is what the typical Nagios page looks like for our signer. Tons and tons of stuff. That's just mostly recapping what I just said. Every single check is for a reason: to address issues that others have had in the past. It's important to note that when DNSSEC breaks, I might have 20 things page me. I want 20 things to page me. I don't want it all "Oh, it was dependent on this thing and that thing was fine." It can be quite [inaudible] when that starts going off. Next.

So, of course, on DSC side – or secondary DNS server side – we're running DSC like many people. We like DSCng as a presenter, although I've seen a couple presentations that indicate there are other DSC frontends available that we may try. Next.

And again, a shout-out to RIPE who provides DNSMON. We love it. That's the only way we can gain visibility into our Anycast network that we're building now. Green equals good.

Questions?

DAN YORK: Thank you. I see a question over here from Roy.

ROY ARENDS: Roy Arends, Nominet. Another question with an additional thing on top of DSC, which is [Hedgehog] developed by [John Dickenson] of [inaudible] and available on the ICANN website. Thank you.

DAN YORK: We have a question in the chat room.

UNIDENTIFIED FEMALE: This was a question back during Roy's presentation. I'll read it off first and then I'll scroll down to where he gives his name.

It says, "Thank you, Roy. Nice presentation. Question: IP addresses could contain privacy-sensitive information. Has Nominet considered the aspects of that?"

This is from Marco Davids from SIDN.

ROY ARENDS: Hi, Marco. This is Roy. So the question was has Nominet considered PII privacy identifiable information. Yes, we've considered that. I'm not a lawyer, so I'm not going to comment.

I don't know exactly how things work in the U.S. in terms of PII. Most of the addresses that we see are resolvers and large resolvers. The specific address that I showed on the screen were infections of individual

machines. I don't think machines are humans. I don't think you can easily personally identify someone by just looking at an address. You need a little bit more context.

But I'm not a lawyer. I also don't know about how long in your specific country you can keep data – if you're actually keeping data or if you actually do this kind of analysis how long you must keep that data. I have no idea.

The way I see it is we've built a car and you can drive the car through a red light or you can kill someone with that car. It's not the car's fault. It's the guy or the girl who's driving it. Hopefully this helps a bit. Thank you.

DAN YORK: Uh-oh, where we could go on that one. So we do have maybe time for one more question if anybody's got one out here for – Sebastien. Okay, go ahead.

[SEBASTIAN]: Sebastian [inaudible] Brazilian Registry Services. So Jacob, how do you solve the problem – well, not the problem, but the situation – associated with the [inaudible] in the signatures for comparison?

JACOB ZACK: We use [inaudible] DNS [read zone] I believe it is to strip out the signatures so that we can compare the non-signature data. Does that sound right?

[SEBASTIAN]: So what if one of your implementations go bananas and the expiration date is wrong or the inception date is wrong?

JACOB ZACK: That would be caught at level two validation via [Valid DNS] which is actually checking – another step of level two validation which is checking the signature validity and expiry times.

[SEBASTIAN]: Okay. So we use OpenDNSSEC as main engine, but we have BIND as our backup. When I was doing that work, I found out [you either] fix the [detail] or use no [detail] or use some general threshold for that. So we made a change and we have some threshold. So [inaudible] within range of each other, they are good to go.

DAN YORK: Thank you. And thank you very much. Let's give a round of applause for all our presenters up here.

[applause]

And these presenters, if you guys could go back to your seats there. And if we need the next DANE and e-mail folks. While we're waiting for them to get up here, we will also put in the plug that you, too, could be sitting up here providing a presentation or case study. We will be going out soon with the call for proposals for the next session in Marrakesh for ICANN 52. So if you're going to be there, why not speak about

something to the greater community? How much longer do I have to fill time? All right, we're good.

Somewhat amusingly, just last night when I was looking at my e-mail I got an e-mail from somebody in Germany who was complimenting me on a couple of blog posts I had written about DNSSEC, etc., but was mostly talking about how they've gone and implemented DNSSEC and DANE in their system service.

Like I said at the beginning of this session, it's been interesting because SMTP has turned out to be one of the interesting drivers for DANE and for the usage of that for people doing that. There are e-mail services in Germany that are advertising on their webpages that they are more secure than others because they provide DNSSEC and DANE support. It's interesting to see where things go with this.

But here to talk about not that, but rather some different techniques to do it, we've got a range of folks. Are you Lynch?

LYNCH DAVIS: I'm Lynch.

DAN YORK: All right, great. First up, we're going to go with Lynch Davis from Verisign.

LYNCH DAVIS: My name is Lynch Davis. I'm with Verisign and I'm just here to discuss some of the work that Eric Osterweil and I have been doing on trying to implement DANE/SMIME on a simple e-mail client.

The prototype goals were pretty simple. We wanted to see if we could actually get an implementation using some sort of standard mail user agent. The implementation that we chose to proceed with was – the current draft is the [-07] draft, but we also selected to include some proposed enhancements, especially the underscore sign and underscore [inaudible] proposed enhancements. Those actually make the deterministic management of the keys from a user standpoint a little easier.

Support of the NAPTR record. I put in the SHA224 encoding, which has become the standard. In previous work I had been doing, I had been using the BASE32. I think that was version three or four of the draft. Next slide, next slide.

One of the things – we did look at a couple different platforms to use, and we ended up targeting Thunderbird for both its cross platform and extension support, had looked at doing more of an integration but chose to try and exploit the extensions framework to avoid recompiling, which was quite an endeavor.

Basically what we've built is, on the right-hand side you can see a breakdown of the architecture. We chose to use OpenSSL and getDNS frameworks for encapsulating most of the functionality that was needed inside what I'm calling libDane, for lack of imagination in naming.

We chose to go ahead and use the C/C++ bindings. The shared library concept integrated very nicely into some of what the Thunderbird framework permitted. The libDANE is completely decoupled from the UI, so if somebody did want to take this framework and apply it in another UI that they wanted to – at least the bulk of the work would be there. Next slide, next slide.

What we had here, since they weren't permitting a whole lot of time for demos, these are some of the screen captures. On the encryption and sending side, the left-hand side would actually be a sending buffer. You can see up at the top there are two miniscule little buttons there. One is for sending. One is for encrypting.

The actual encryption itself takes place by capturing the actual written buffer and it passes it off to libDane. libDane then will use actually the SSH224 from OpenSSL to hash the senders address, and then getDNS will actually retrieve the valid key from DNS. Again, OpenSSL does the encryption for us and passes back the encrypted buffer and pushes it into the buffer for sending. Next.

On the decryption side, we did manage to find a hook that when the user selects an e-mail message it looks for the actual signature saying that it was signed the encryption header, it will decrypt it and replace the buffer with the decryption.

Luckily, the way this works is that decryption is only done in memory. So it avoids writing any files and it preserves the security of the file. Next screen.

As a side benefit of the way that worked, if you did try to forward or reply to the e-mail, Thunderbird actually goes back to the original text of the e-mail in memory, which actually is still encrypted. So if somebody wants to forward on the decrypted, they're going to have to go a little bit out of their way to do the cut-and-paste. No way to prevent that at this point.

And then on the signing message verification side, again we have the message signature validation button. So again it will take the text, and it's the same process. It hands off to libDANE, libDANE again for signing is going to look for the DNS entry with the underscore sign attribute with the entry and it comes back with a fully-signed buffer.

When it's received, it's more of an on-demand type of thing, but it will verify whether the signer's there, whether the signature exists and whether it's a valid signature. Next.

Some of the things that – you'll see on the right-hand side this what a public key record looks like inside of our Verisign's managed DNS entry. We did run into some of the challenges. Finding a provisioning platform to provision these records – if somebody isn't up-to-date on at least a little bit technical, dealing with certs continues to be a challenge just from the nature of certs. But to provide the proper encoding and get it into the record manually is quite a challenge. I found it challenging at least.

So what we're missing there is tools for enterprise. Enterprise would obviously need to automate this type whether they're currently using existing AD or whether they're starting from the ground up. And for

individual users, the registrars have got to be able to step up and provide some tools for the folks that are using this.

Again, dealing with the actual certs, any way to abstract that from the end users or whoever is maintaining these is definitely going to help.

Where the libraries are currently, lacking configuration in the locations of the certs are fairly much hardcoded at this point. We did run into some interesting issues with mail changing over the years. I guess there's [format equals flowed] for content is a parameter for mail. It allows mail to be visible on various devices and reformatted on the fly. Well, it kind of breaks signing when your mail user agent is inserting and deleting white space on you as you go.

The other things – for the prototyping that [inaudible] been dealing with mostly with self-signed certs rather than certs that have a true signing certificate authority, need enhancements to allow for the private keys and the keys stored locally to accept passwords and passcodes.

The last one is irrelevant. I managed to solve that between when I submitted the slides and now.

UNIDENTIFIED MALE: Cool.

LYNCH DAVIS: That's our contact information, and that's kind of where we're at.

DAN YORK: Very cool. So this is your implementation of this [inaudible]. Any questions from folks here? Oh, come on. Somebody's got to have a question here. Any questions remote?

You're not going to let Eric and Lynch off this easy, guys. All right, well, we'll see if we have more questions when they come back with this. Thank you. Let's go on for the next one.

I should note, too, there's this draft. The draft is still under discussion in the – which group is it in? DNSOP or DANE? It's in DANE on the DANE side. Who's involved in that draft? Oh, Paul [Hoffman] is back there. I see you. Do you have any questions?

So the draft is out there for discussion. People are encouraged to read that draft and provide discussion to the DANE mailing list as well. Eric wants to jump in. All right.

ERIC OSTERWEIL: I just wanted to get my voice out there. I wanted to hear myself talk. No, real quick. I just want to say there's an active discussion going on. A lot of the things that Lynch learned and got through with the really good implementation work he's done really helped set some of our pains and postures on this.

So if people are interested in being able to download a plugin where you can click a button and automatically encrypt your e-mail to someone else, another organization, we definitely encourage participation.

DAN YORK: Jacques?

JACQUES LATOUR: Jacques Latour, .CA. I'm on the program committee. Do you think by the next ICANN meeting you could do a live demo of this?

ERIC OSTERWEIL: I'll speak for Lynch, even though Lynch is the one who's doing all the real work on this one. I think we actually just had a miscommunication. I think we probably could've done one today. We can definitely do that. I think we're planning to do a round of this sort of stuff. I don't want to speak for working group chairs, but maybe even at the IETF.

DAN YORK: Cool. We're always interested in demos with the subject, but that's always a little bit challenging sometimes, but that's cool. Great. And can people get the plugin and use it?

ERIC OSTERWEIL: We're working on making the plugin available, but I've got a little more cleanup to do and we will eventually be positing it out on GitHub.

DAN YORK: Cool. Let us know. We'd love to promote that, I'm sure.

ERIC OSTERWEIL: It's definitely in the plan. There's no [inaudible] about whether we'll do it. It's just when.

DAN YORK: Thank you, guys. Next up Wes Hardaker is coming back to talk a bit about some of the work that he's been doing on this.

WES HARDAKER: I'm Wes Hardaker from Parsons again. I'll be talking today about the server-to-server level of DNSSEC, the hidden. Next slide.

I'm going to talk about few things. I'm going to show you quickly how it works. If you've forgotten [inaudible], I'll talk about some vulnerabilities and why you need to secure SMTP traffic from server to server, and I'll talk about how DANE and SMTP working together really come to the rescue and it's really the only forward-thinking solution going forward.

Then finally, I'll talk a little bit about where we are in terms of implementation and deployment, because there actually is some. Next.

Really briefly, if we have Alice who needs to send a mail to Bob and Bob's far away and Alice needs to send mail from her ISP to another ISP if they have different suffixes to the right of the @ sign, what really happens is Alice isn't talking directly to Bob's ISP. Alice is talking to hers and to the mail transfer agent at her ISP and says, "Here, I have this mail for Bob."

Her SMTP at her ISP says, "Oh, I can go find out where Bob is," and sends it to the mail transfer agent at Bob and says, "I have mail for Bob. Please deliver this for me." And then eventually, Bob uses IMAP or POP or something like that to download it. Next slide.

The piece that I'm talking about today is the middle piece. I'm not talking about how Alice sends the mail and I'm not talking about how Bob receives the mail. I'm talking about how the two servers communicate between each other, although there is some similarity, especially on the left-hand side of the diagram. I'm really concentrating on the middle. Next.

It works sort of like this. Alice's ISP says, "Well, where should I send this mail to bobsisp.com?" and asks a DNS server – Bob's DNS server. Well, Bob's DNS server would respond with, "Well, you should send it to mail.bobsisp.com and the address for that is _____." They actually connect to the mail server and says, "Hey, I've got mail for Bob." Next.

I really wish it was that simple. It's actually much, much worse than that and we'll go into a little bit of the details, but I can't do it all because it's much more than I can do in a short period of time.

There can be multiple DNS servers, for one thing. Every domain really should have at least two. Alice's mail server actually doesn't ask Bob's DNS server directly. Alice's mail server probably asks their internal ISP, so there's a whole other resolver involved that wasn't in that last diagram. And there's probably multiple resolvers, too. There can also be multiple mail servers for both sides. Go ahead and next.

So in the reality – and I'm not going to read through this one. And this isn't even fully laid-out, but this is closer to the truth in terms of how much traffic actually goes on to send one piece of e-mail from one server to the next. Again, I'm leaving out some stuff because I'm not doing all the DNS servers and all the MTAs out there. Next.

So this is the exact same slide. I wish it was so simple. Those are the types of things that can go wrong. Next.

So what can go wrong? First off, there can be multiple DNS servers. Some of them could be compromised. That's always a possibility. Alice's mail server asks her ISPs resolver. That resolver could be compromised and have incorrect data in it. Compromised not necessarily completely. Compromised could mean cache poisoned in this case. There can be multiple mail servers on the other side, too, and they could be compromised.

Then, finally, the one that was missing before is you could have a man in the middle. Somewhere in the middle is somebody intercepting the mail and either keeping it or even pushing it forward. that's always my biggest scare because it's fairly easy to be a mailman in the middle and have nobody notice. It's too easy. Nobody looks at all the headers to see where it went through and things like that. Next.

So there's a solution, fortunately, and that solution is actually very married to DNSSEC. So DNSSEC sort of fixes the first two compromised problems because the data can't be compromised even if the sever can and you can detect when the data has been modified. That's the purpose of DNSSEC, to detect when the data has been modified.

So even if the DNS servers are compromised, they can't tell you bad data. And then most recently, the use of DANE and DNSSEC together provides you protection against the man in the middle himself. Next.

So let's talk about some of the [SMTP] vulnerabilities. Where is this coming from? First off, if you have a domain name that is going to

receive mail, the way you advertise your mail server is you advertise an MX record. That MX record could be spoofed. There's an A record or a quad-A record associated with that as well for the address for the mail server.

If anybody spoofs that, you're toast. They can convince you to send mail for Bob's ISP to Evil Eve somewhere off in the corner, and there's no way around that if they control the DNS. So the trick is don't let them control the DNS and that's what DNSSEC is about. It protects those records so that you can tell if they have been modified.

And then finally, eavesdropping is quite easy. SMTP is unencrypted by default. In fact, most of the world's traffic up until somewhat recently has been entirely unencrypted. That's hard. Most recently – there's a couple of major mail service providers that have started encrypting stuff between each other at least, and that's actually improved things a lot.

And then opportunistic encryption helps, which is "I'm going to go talk to Bob Server and I'm just going to start encryption. I don't actually know if I'm talking to the right Bob Server, but at least it's encrypted so that nobody in the middle can watch it if they aren't actually participating in the conversation." However, you may just be encrypting it to possibly the wrong person. Next.

So if DNS actually gets spoofed – in other words, if I tell you that example.com's mail server is actually EvilBob.com, you get a main in the middle because I'll try and go connect to Evil Bob's server instead.

Again, SMTP is unencrypted by default, but it's also unauthenticated. You have no notion if you're talking to the right mail server or not. Even

if you have the right address, there can be a man in the middle that's sitting on the right address, but catching the traffic first.

So opportunistic encryption helps but it's not perfect because you really need the authentication to go along with it.

What's even worse is that because SMTP – in order to turn on authentication and encryption, which means turning on the TLS layer – both sides have to say, "I do security." And a man in the middle can sit there and say, "I don't do security," and you have no choice but to believe him because you have no idea whether he actually is supposed to do security or not.

An evil man in the middle, even if you were supposed to be doing authentication would just say, "I don't do security," and you continue talking to him unencrypted and unauthenticated. A [CA] solution really doesn't help in this case because of that. The man in the middle just says, "I don't do security." Not only that, if the DNS is compromised they can send you to whatever name they want, including one that they've properly registered under a [CA]. next.

So this is where DNSSEC and DANE win. DNSSEC and DANE solves all of these problems, because again, with DNSSEC you can believe that the MX record, which is the mail server record, that led you to where you're trying to connect to is correct. It's guaranteed to not be modified unless they have much worse problems.

The TLSA record that is pointing you to their certificate to do the authentication and the encryption with to start that process at least is

verified that you know that you're getting the right TLSA record and it hasn't been modified.

And then with the TLSA record that you now know is good, you know that the certificate that you're receiving from the other side is actually the absolute correct one. They present a certificate and the DANE TLSA record says, "This is my certificate. You should make sure that it matches," or possibly "This is my [CA] and please don't accept anything else you ever receive from anybody. If it doesn't match, you kill the connection."

The other thing that it does – the TLSA record, the way that – and I'm a coauthor of the draft, along with Viktor Dukhovni – we've defined it so that if a TLSA record exists, you must expect to do security. So if somebody says, "I don't do security. I don't support the start tls command," you can say, "Nope, I'm not talking to you." And you can believe that. So it verifies that you've connected to the right place. Next.

I'll talk a little bit about deployment. The DANE TLSA support is now in Postfix 2.11. It's out there. It's the primary release, the most recent release of their software. It was released in January, I believe. And on the server side, doing configuration is fairly simple. You just publish a TLSA record and you give it the path to the keys and off it goes.

On the client side, if you wanted to check to see if somebody is using DANE and DNSSEC, there's only two really configuration settings that you must use.

The one thing I will say is that it does use the local resolver or it does use whatever resolver. It's not actually doing DNSSEC checks inside the

application itself. It's expecting you to be talking to a resolver you trust along a trusted path i.e. the local host. So the recommendation is that you have a resolving, validating name server on the local host and it looks for the AD bit and it believes it, so do take that with that understanding.

The [EXIM] implementation is underway and it's expected to be completed in 2015, which is the second-most popular mail server out there these days. So those are the two biggest ones.

The known large early adopters, there's actually quite a list. You can see that there's some pretty big ones on there including mailbox.org. A bunch of work has been undertaken in Germany to get most of their mail servers using DANE and DNSSEC. The IETF, debian, FreeBSD.org. I think there's one other that I've now forgotten off the top of my head.

The important thing is there's actually 270 of these out there already turned on and using it. That's huge. The reason that's huge is the specification isn't even published in an RFC yet. It's soon. We're working on it. It's coming soon. But 270 people are using it in advance of it being out. That's far, far faster deployment than DNSSEC itself was seeing. Next slide.

Questions? I think you'll recognize the picture. Do feel free to catch me if you have any questions or you want to discuss things Viktor Dukhovni is actually the primary offer of the Postfix implementation and I have been conversing with him for now years, I think, so I can probably even answer some of those questions or contact him if you have questions about that.

DAN YORK: Thank you, Wes. This is great stuff. We do have time for a couple questions. I note, too, that the man in the middle attack that Wes was talking about here, back in September, in early September – September 10th – a team at the Carnegie Mellon [Cert CC] released a study. You can find it on their blog. I wrote about it on the Deploy360 if you can't find it there.

They documented that they found in their research that there were people doing this, that they were poisoning. They don't know who. They identified a list of IP addresses that they had found that records were being redirected to. But they were finding that a large quantity of e-mail in some parts was being redirected the MX records were being poisoned and redirected to some servers. Presumably it was being delivered through, or else people just missed the e-mail or whatever.

This is a very real thing that people are seeing out there. It's not an abstract attack. This is very real. I'd encourage you to look up that work from September 10th but the [Cert CC]. Any questions for Wes?

ERIC OSTERWEIL: Hey, Wes. That was a really good presentation and really good work. I think it's awesome.

WES HARDAKER: Thank you.

ERIC OSTERWEIL: I have one question. It's not thinly veiled. It's just an honest question. Have you guys thought about looking at whether getDNS is useful to you guys as far as the local resolver [inaudible] and it being able to do full validation?

Just if that would help get over that—

WES HARDAKER: Yeah. So first off, I didn't do the implementation in Postfix. That was entirely Viktor. And Viktor and I disagree a little bit on where validation should be done. To give his side of the story, because he did it, he does not believe that he wants to add the complexity of doing validation inside of an application to the application.

So in his view, he would much rather just for somebody to use a local resolver on a local host and he tells people how to do that and stuff. He does not want to add that complexity to the application.

There's very different trains of thoughts. I think I mentioned in one of my previous talks, I'm more for putting it directly inside the application. But he wrote the code, not me.

DAN YORK: That's one of those that can devolve into a religious war on some level. Any other questions?

Wes did mention that there is a draft – or a couple of drafts – that Wes has contributed to around these topics within the IETF DANE Working Group and I would encourage you to take a look at those and send feedback and comments to Wes if there are any.

Okay. Well, then let's bring up our last presentation. Scott Rose from NIST is back to talk about some of the work they've been doing with some other different spins on this.

SCOTT ROSE: Hi, this is Scott Rose again. I'm just going to do a brief overview of what we're working on right now at NIST as part of the what we call the high-assurance domain project, which is supposed to be a joint project between NIST, Homeland Security, and the Financial Services Coordinating Council which is actually part of DHS. Next slide.

First, going back, I mentioned the Tiger Team that was formed in 2011 when it was trying to do DNSSEC actually had a second goal, and it was authenticated e-mail or trustworthy e-mail.

The ultimate goal was to have everything, what we call government-to-government. Anything between government agencies as well as anything from government to a citizen would be authenticated, or at least be able to be authenticated.

Because right now if you go to places like the irs.gov website, they have an explicit statement saying, "The irs.gov will never send you e-mail." If you get an unsolicited e-mail from irs.gov, just delete it because it's spam or a phishing attack.

The idea was that maybe one day wouldn't it be nice to change that so you can actually get some assurance that it actually came from the IRS?

So they decided to try and do kind of a walk before we run thing, was deploy things that are already out there not that aren't really part of

DANE, like the [Center] Policy Framework, [Domain Keys], later DMARC. The Tiger team was actually before DMARC was actually announced. So those were kind of the goals, was trying to get those deployed, at least within .gov and try and get – to encourage people to actually use it.

And DNSSEC was seen as the enabling technology. If you're putting policy and important stuff in the DNS, you should probably secure it using DNSSEC and to check those DNS signatures. Next slide.

So now we get to what we're doing now. We kind of see it as two different levels. One is looking at the enterprise, like mail going between from one enterprise to another, and as well as end to end, from one user to another user. So we're lumping – I don't know if we're right or not – thinks like SPF, DKIM, DMARC and as well as SMTP over TLS as kind of enterprise-to-enterprise things. These are like what MTAs are going talking to other MTAs.

As well as we're doing the end-to-end stuff, which is user-to-user, and that's kind of the two ones that we're focusing on are OpenPGP and S/MIME. So, last slide.

This is kind of what we're doing. If you go, the bottom there is kind of our project page. You'll notice it's .com because we thought it was going to be a joint project, so it wasn't going to be in a .gov space. But it's ours.

We find the usual NIST networking goodies. We have a monitoring for these kind of security artifacts both in gov, list of banks and the edu. We're looking for SPF, DMARC records. We're trying to look for DKIM.

We're basically trying to guess. If we haven't seen an e-mail, it's kind of hard to know.

We're developing. We have interactive test tools for SPF, DKIM, DMARC and OpenPGP. S/MIME A is still kind of under development. We have it. The guy who wrote the original tool left. He was a guest researcher. He left to go back to his original job. We still have the code and we're going to be putting it up online soon.

The next big thing is to have a guidance doc. Everybody may know the special pub series, the 800 series. We're going to be doing a new one on trustworthy e-mail as a service, which if you look at other NIST guides, they never actually talk about services. They always talk about servers. So if you look at the e-mail guide, it's all "how do you secure an SMTP server?" This is going to be how would you look at securing e-mail as an entire service that you're providing.

As well as contributing to protocol specifications in the IETF. So it's kind of what NIST is doing. We encourage obviously people to go to our website, send us questions, comments, gripes. If you want to get involved, please do. We'll be announcing the special pub draft hopefully sometime next year. So we're looking for people – especially e-mail experts, if absolutely knows anybody – to help comment on that, because we're not but we're writing it.

That's kind of what we're doing. That's our state of affairs from within the government as we see the attractiveness of using the DNS to support trustworthy e-mail.

DAN YORK: Awesome. Thank you, Scott. So if people do want to get involved, I know you mentioned you have all these tools. They can just go there and use them, right?

SCOTT ROSE: Yes. The interactive tools are online. You can interact with them. They're pretty basic. You send an e-mail, you get an e-mail back, that sort of thing. The OpenPGP key tool is up and active, but if you don't have any OpenPGP key resource records in your zone, it doesn't work so well. You get it back in the e-mail message saying, "I couldn't find your key." But if you are, it does work.

Like I said, the S/MIME A is coming soon.

DAN YORK: We should note there is a draft out on the OpenPGP element of this that Paul Wouters, who Wes was channeling earlier, is involved in that is also in one of the groups out there that you can go and take a look at and comment on.

Any questions for Scott?

We've been sitting in this room for six hours talking about DNSSEC and we're tired. Scott gets the benefit of being the last presenter.

Any final comments from the folks here?

Okay, great. Well, then let me just bring up our last set of slides as we wrap this up. I'll wait for the people who are remote. Okay.

We always like to just end this for people who are out there, if you are a TLD operator, registry, any of those, we ask you to please sign your TLD, [accept] your DS records. We have a thing out there. Work with your registrars.

One other element at the end. We talked a number of times about metrics. We increasingly would like to ask – we'd love to see some better metrics.

We mentioned at the big the [hold] issue around the signing – and this is a particular one for TLD operators. We have these high-level charts of how many TLDs have been signed, but we're really looking to go to that next level, the second-level domains. And there are some folks who have done that. Verisign has done that with the com, net and other places there. Some of the other different domains, the folks at CZ, NL, SE have provided various different charts.

But we are alga ways to try to dig deeper into that, so we're looking for help with that. So any degree that statistics can be exposed would be awesome in some level. Next slide.

We of course are asking zone operators to go and again sign zones, talk to registrars. One of the questions we keep getting from registrars, even with the 2013 RAA coming in there and stuff is a lot of registrars keep saying, "Nobody's asking for DNSSEC." So one of the big things we really always encourage out of these meetings is for everybody who's here and everybody who's listening or listens to this remotely, please ask your registrar if they will support DNSSEC, at the very least to accept DS records. Big deal on here. Next, please.

Network service providers – again, operators, ISPs, anyone out there, enterprise networks – we've heard a couple times here how easy it is to get validation running in those networks. We just need to get that working out there. Also, we're encouraging people again to look at DANE and how we can support that, in what ways. Next slide.

For website content providers, we again are asking people, please, sign your zones. Let's talk about DANE. Let's look at what's out there and look at how we can get again more DNSSEC validating resolvers. We'd like to see Geoff's chart go up even more. We're already at 12% as of—

GEOFF HUSTON: 13.

DAN YORK: Oh, 13. We hit 13%. Okay. Good deals. Let's keep going up. That's good. Next slide.

And again, everyone, we would ask you to please do this. Share your lessons learned. Again, as we mentioned – I'll put another pitch in – we have another one of these sessions coming up at ICANN 52 in Marrakesh in February and we would encourage you to think about what kind of case study, what kind of tool, what kind of service, what could you present to this group to help spread the knowledge about what we can do collectively to make DNSSEC work better and to bring about a more secure Internet? We would ask you to do that.

I want to say a special thank you to all of our presenters and to our participants. I'd also again like to thank Julie, too. Let's give a round of applause to everybody who's been involved here today.

[applause]

Oh, there's a comment in [inaudible]. While we're getting that, I want to also mention, too, at the beginning when we were talking about the DNSSEC deployment maps and things, I mentioned – I showed the map of Africa in the range of ccTLDs that weren't signed there yet or an area of growth.

It was kind of interesting. As I was looking at Twitter moments later, I saw that there is a post up on the ICANN blog about their efforts that they're continuing to focus on around helping to do that. So good to see that kind of work happening out there. I know a number of folks are working on DNSSEC programs within Africa in particular.

We have a comment from the remote room.

UNIDENTIFIED FEMALE: Yeah. So this is a comment from [Ed Louis]. He says: "Suggest a set of stats, a sample report when you're asking for statistics." I think he means if you wanted to show the stats that you were looking for in a sample report.

DAN YORK: Thank you, Ed. Sure, I'll be glad to suggest [inaudible] what I'd like. I'll be glad to do that. [laughs]

To go on that point, too to all of this, it's been great to have Geoff's charts. We're always looking for more kinds of ways that we can show what is the true state of DNSSEC deployment and measurement.

For a long time on that particular one I'm mentioning, we've been showing this TLD – the number of signed TLDs. That's awesome. But now we're at a point where most TLDs in the grand scheme of things, a lot of the TLDs are signed. Yes, we've got ccTLDs still to work on, but a lot of the major ones are out there. Now we need to go to that next level and see what's really happening at the second level and how we can start to move that.

Geoff wants to say something.

GEOFF HUSTON: Do you want the really good news or the not-so-good news? One quarter of the world's users refer their queries to DNSSEC-validating resolvers. Half of them hate SERVFAIL responses, and then go and use a dud resolver because they just can't take no for an answer.

So if we get rid of the duds, you'd double overnight. Isn't that amazing? One-quarter of the world actually [inaudible] DNS Key queries.

DAN YORK: Very cool.

GEOFF HUSTON: v6? 3%.

[laughter]

DAN YORK: Ah! Hey, I think it's 4.5 now if you look at some of the stats.

So let me just say thank you again. There are some websites up here: dnssec-deployment.org, [inaudible], and dnssec-tools.org. Actually, we're missing [inaudible].

Thank you all for participating. We will see you again in Marrakesh. And please feel free to stay around for a few minutes and have conversations. We do need to actually exit the room, though, because there is another group coming in. But thank you all for being here. See you in Marrakesh!

[applause]

[END OF TRANSCRIPTION]