DURBAN – DNSSEC for Everybody
Monday, July 15, 2013 – 17:00 to 18:30
ICANN – Durban, South Africa

DAN YORK: Good afternoon. We'll get going in just a moment. Please take a seat. If you did not get one of the handouts that we have there are some on the table at the back there. So please go and grab one if you're here. This is the DNSSEC for Everyone Session. If you're looking for something else you're in the wrong place.

We're just taking an extra moment to configure the sharing into the room because we do have a number of remote participants coming in to watch this. So in just a moment we'll get going.

JULIE HEDLUND: Thank you everyone for being patient. I'm happy to see some more people coming into the room. My name is Julie Hedlund, I'm with ICANN Staff and we're going to get ready to start the DNSSEC for Everybody – A Beginner's Guide Session in just a moment. Thank you.

DAN YORK: Good afternoon. Welcome to the DNSSEC for Everybody – A Beginner's Guide Session. We are going to bring you an interesting afternoon. We're going to talk a bit about what DNSSEC is, how it works, give you a little skit, a little performance of how DNS works and pieces like that. We will have time for questions and answers.

We're using the microphones up here for our skit but when we go to questions and answers we will put them here in the stands and we'll ask you to come to the microphones and speak your questions, because we do have remote attendees who are joining us here.

So welcome. My name is Dan York, I'm with the Internet Society and part of the cast of characters that are working with us today. We did pass out to those of you who are in this room this sheet on session notes for DNSSEC for beginners. It shows an overview of the schedule that we're doing and it also has a list of resources on the back.

If you did not get one we do have some on the table at the back there that you're welcome to pick up. And it's also available from the website where you can go and download that as session notes. The presentation is also available on the website so you can see the material we have here. Obviously the skit will be live for your entertainment value here.

Let me introduce our cast of characters that we have here: Russ Mundy is seated right here. He'll be part of what we talk about today and Russ will also be speaking a bit about some examples of DNSSEC when we get to that a little later. Roy Arends is here from the UK and he will be an actor in this as well. Just as a note your schedule says that Roy is doing a section but you'll see that I am going to do that. Roy is just experiencing some voice challenges right now.

We have Norm Ritchie, who is also seated here. And we have Jacque Latour. And we have a special character who is out in the audience who will also be assisting us in our skit; our mysterious Dr. Evil will appear at some point in time.

Okay. So with that we will get going. Let me begin with the presentation. How many of you are familiar with DNSSEC? Yes, I see a few ringers in the crowd. How many of you have no idea what DNSSEC is whatsoever? Okay, good. How many of you are familiar with DNS? Okay, good. How many of you have no idea why you're here? [laughter] Okay. "They told me in a novice session to come so here I am."

Okay. So we're going to go through this schedule and work with this. We're going to begin with the origins of DNSSEC, all the way back in 5000 BC – that's when they had the Internet, right? At that time we had Aguinna. She lives in a cave on the edge of the Grand Canyon and this is Og. He lives on the other side of the Grand Canyon.

And as you would expect in a story Aguinna and Og want to get together, but it's a long way to go from one side to the other. So on one of the rare visits they notice they have smoke coming from Og's fire so they figure out a system where they could signal each other from one side of the canyon to the other and they can communicate and do all this type of thing. And it works great, they have their conversations and all is good.

Well, one day the mischievous caveman, Kaminsky, moves in next door to Og and he starts sending smoke signals as well and all of a sudden poor Aguinna can't really tell which one is which. She sees these signals coming up but she has no idea what they mean. So she sets off to go down the Canyon and get to the other side and try to figure out what's going on. She sees these two signals and doesn't know which one is the one she wants.

So they go and consult the village elders. And caveman Diffy has an idea about how to go and do this. So he runs to the back of Og's cave and there he finds this pile of strangely colored sand that has only been found in Og's particular cave. What he does is he runs out and throws that into the fire and suddenly the smoke turns blue.

Well, now, because this is only found in Og's cave, suddenly they're able to have their conversations again; Aguinna and Og can chat again because now she knows that the blue smoke is the one that she has to pay attention to. She has to look for the blue smoke. And that, basically, is what we're going to be talking about – how do we get the blue smoke? How do we identify that this particular address, this particular piece is the one that we need to be paying attention to?

Onto a little bit about DNS and the pieces work in here. If we think about the high level of DNSSEC we have, at the root level, the various top-level domains, we have the various sub-domains that are beneath that and so on and so forth. We have this whole structure that's in here. At a base level, when you go and ask for an address your local resolver that could be running on your laptop, your mobile device, it could be running on your PC, whatever – somewhere you have a DNS resolver.

That piece of software goes out there and it traverses through the DNS hierarchy looking for whoever can give the correct answer, and it goes on and goes on and does this. Once it has an answer that resolver, that software running on your computer or mobile device caches, holds onto that answer for a certain period of time. So then if you want to go there again you can get the answer back quickly.

So when you go to ICANN.org your resolver goes out, finds that information, gets it and holds onto it. This is how DNS works. That's Dr. Evil who's out there. Look at that, he's tearing things apart! We'd better go and find him. So to illustrate this [laughs] we're going to bring up our cast of characters to talk a bit about this. And we all have… Excuse us for a moment while we get some taped things on us.

Norm is going to do the Joe User and he's also going to do the talking through things bit. We should mention that Norm went through some great work to create some brand new t-shirts. For those of you who've seen us do this play before, if you ever saw it before, we had really aged t-shirts.

But Norm went through the work to go and create some great new t-shirts, awesome things, and they are somewhere between Cairo and Johannesburg right now because Norm is here and his luggage is not. All right, are we ready? Okay.

JOE USER: Okay. So we're going to do this skit in four parts and we're going to show the different concepts of DNS. The first thing we're going to do is just a simple DNS query. Dan just explained how DNS works, well, we're going to act it out. So up here I am Joe User; I'm your typical person that's going to be using the Internet. We have my ISP that I pay my monthly fee too who is my resolver DNS. And the root DNS servers, comp DNS servers and in this case, big bank.

So here's an illustration, very slowly, of how DNS works. I'm at home and I go and do some banking. I do a lot of banking, I have a lot of bills.

So I go to my laptop and I type in www.bigbank.com. That gets handed off to my ISP.

ISP: Thank you. I am an ISP and I don't know a lot. I don't know anything. So I've got to figure out where bigbank.com is. The first thing I do is I go to the root and I ask, "Root, where's bigbank.com?"

ROOT: Hmm, I don't know where bigbank.com is but you can find out by going to the .com name server. And is that 1.1.1.1.

ISP: Thank you very much. Mr. .com, do you know where bigbank.com is?

.COM: I don't know where www.bigbank.com is, but I do know where bigbank.com is and you can find bigbank.com at 2.2.2.2.

ISP: Thank you, I'll go and ask bigbank.com. Do you know where www.bigbank.com is?

BIGBANK.COM: Yes! I do happen to know where www.bigbank.com is. It is at 2.2.2.3.

ISP: Thank you. Hey Joe. They don't do v6, only v4. The address is 2.2.2.3.

JOE USER: Thank you Mr. ISP. Now I have the address for www.bigbank.com and my computer can go off and find the bank and I can do my banking. So what you saw there, I'm the user, I didn't have to do much other than ask the question, wait for my recursive resolver to give me the answer and he's spent all this time talking to the DNS authorities who had the actual answers. They did all the work I just sit there, ask a certain question and wait for the answer.

This obviously happens very, very fast on the Internet. This is a bit slower and now we're going to redo the skit. And this time we're going to show you what was a flaw in the DNS – it's inherently insecure. As you can see, they were just passing information across and waiting for the answer. Now we're going to show you what can happen. It's called the "man in the middle" attack.

Act II. More banking to do. Same scenario. I sit down at my laptop and go click-click: www.bigbank.com and hand it to Mr. ISP.

ISP: You again? www.bigbank.com. I don't know where it is so I'll go and ask the root. Do you know where www.bigbank.com is?

ROOT: I don't know where www.bigbank.com is but I do know who has .com. Go see the name server at 1.1.1.1.

ISP: Thank you. 1.1.1.1. I'll go and ask Mr. .com. Do you know where www.bigbank.com is?

.COM: I don't know where www.bigbank.com is, but I do have information about where bigbank.com is and you can find it at 2.2.2.2.

ISP: Thank you. I'll go to 2.2.2.2 and I'll ask them. Mr. bigbank.com, do you know where www.bigbank.com is?

DR. EVIL: Hello, I have an answer for you. www.bigbank.com is at 6.6.6.6.

ISP: Thank you very much.

DR. EVIL: You're very welcome. Any time.

ISP: Hey Joe. The address is 6.6.6.6. Have fun.

JOE USER: Thank you very much Mr. ISP. Now I can go off and do my banking at my favorite banking spot at 6.6.6.6. And my computer is going right to Dr.

Evil, knowing no difference. So that's called the "man in the middle" attack. Now we'll introduce DNSSEC, we'll introduce the blue smoke. And the way that's done is through what's called… There's a signing and some details you probably don't care about, but there is a chain of trust. As you can see, the DNS is set in a hierarchy.

When you use digital signatures though, if they trust each other they can validate each other's response. So we're going to introduce DNSSEC signing. Okay. So what they're doing now is the root is now signed. .com is signed. Root and .com know each other, they exchange credentials. And good old bigbank.com gets signed. Yay. And they're going to exchange credentials. We now have introduced what's called the chain of trust.

So they can validate responses now. So they can actually check when they get a response; they can actually go back and validate that. It came from that server, that person, and it had not been modified in the process. So let's redo the skit again with a signed DNS. More banking. I have a *lot* of bills. I'm going to sit down at my laptop again and go to www.bigbank.com and I'm going to hand it to Mr. ISP.

ISP: You again? Bigbank.com. Okay. I don't know where it is so I'll go to the root. Do you know where www.bigbank.com is?

ROOT: I don't, but I know where the name server for .com is. It's at 1.1.1.1 and here's my signature to tell you that that is in fact the right place to get it.

ISP: Thank you. .com, do you know where www.bigbank.com is?

.COM: I don't know where www.bigbank.com is but I know where bigbank.com is. Here's the signature you can validate and you can trust. And you can find it at address 2.2.2.2.

ISP: Signature looks good. Thank you. Hello bigbank.com. I'm looking for www.bigbank.com.

DR. EVIL: Hello! I have a perfect right answer for you. www.bigbank.com is at address 6.6.6.6. And don't lose time with formalities, here's the answer.

ISP: Let me check this… 6.6.6.6 signature. [groans] It's no good. Oh, here's another answer.

BIGBANK.COM: I do have an authenticated answer for you, properly signed through the DNSSEC chain. It is 2.2.2.3.

ISP: Thank you very much. It's good. Hey Joe, here's the address: 2.2.2.3 and it's verified.

JOE USER: Thank you very much Mr. ISP. Now I go off and do my banking with a signed and authenticated response from the correct bank. So that's how DNSSEC works. An important thing to note here is that I, as a user, don't have to do a lot other than send him my request. The resolver is going to do all the authenticating.

The authorities are going to do all the signing but the average user really doesn't have to know much about DNSSEC other than the fact that it's there and it's more secure. So that's the end of the play. Thank you. [applause]

DAN YORK: All right. Thank you to all of you. I hope that helped explain a little bit about the DNSSEC process. We should also thank Andre here who came in as our Dr. Evil. [laughs] [applause] Just to review a little bit about what we saw there, again, we have this concept in our scenario of Aguinna – the resolver – is chatting with Og – the server – and when she's confused and doesn't know what it is she is looking for that blue smoke.

She's looking for the fact that the resolver can verify that that is the correct answer that Og is sending; that the real Og is sending that message. At a basic high-level of DNS, part of what we saw here is that there really isn't security built in to the DNS as it is. We can have "man in the middle" attacks where someone can go and do that.

The basic reality is that when the resolver sends out the message and says: "I want to know what www.bigbank.com is," which answer they take is whichever answer they get quickest. Speed wins. So as you saw here, both the regular resolver for bigbank.com and our Dr. Evil both sent their responses. What we sort of showed you was that Andre jumped in there the quickest. He was the fastest one to get back to the ISP. So speed wins.

So if an attacker can get himself or herself to a place in a network where they can go and send the packets quickest, they can get in there and do that. Now, the other part of what we didn't show in the skit is the fact that once the ISP had that answer, he would cache it and hold onto it for a certain period of time. So it's not just that he gets the bad answer once, he gets the bad answer and then he holds onto it and he keeps giving it.

So when Joe User kept asking for www.bigbank.com the resolver would keep giving it back until there was a timeout. It's what we call cash poisoning within the security space but it's this idea that he would hold onto it for a certain period of time. So Dr. Evil would say: "Here it is. And hold onto it for a week," and the ISP would go and do that.

So this is the whole concept of DNS. And what DNSSEC is all about is preventing the wrong answer from getting in there. Of providing this idea of blue smoke and this idea that this is the correct answer, this is how it is. And that's what the signatures are about and that's what the resolution is about.

Now, I'm going to ask… Okay, DNSSEC uses these digital signatures. So there is two parts to DNSSEC. One is you saw us going up here and putting stars on our names. We have to sign each of the records that are in there, and you as the owner or bigbank.com or whatever it might be have to sign your domain. That's one side of it. You give it the star, you give it the blue smoke. That's part of what's there.

So you're creating new records; there are some new records that are out there. One called a DNSKEY, another one called an RRSIG. There is some stuff that gets down in there, which we'll talk a little bit about. Those are the pieces there. You create these new records, you sign it, you give your domain that signature, that blue smoke.

The second part of that is that you need to have a resolver that checks. If Jacques was coming along here and he was the resolver and he didn't check signatures, it wouldn't matter. It wouldn't matter if we signed these. You need these two parts to make DNSSEC work: you need to have the signatures and you need to have the resolvers.

Now, the other part is that Dr. Evil could have come in there and he could have tried to say that he had a signed signature. He could have said, "I've got a signed result," but what would have happened then was that the resolver would have looked at that and said: "Well, does that signature match? Can I track it all the way back up to the root?"

There is this concept to what's called the chain of trust when you get into DNSSEC, which says that all the way from the top of DNS down, I can know that the answer that the person operating DNS for bigbank.com, they put the record in there for www. I know that that is

what they wanted me to go to; not something that somebody else said. That's what this is all about.

Is the information that you are getting out of DNS the same information that the person that owns that domain or registered that domain put into DNS? That's it. Can you get out what was put in by the original person who has control of that? That's what this is all about.

And with that I'm going to have Russ Mundy speak a little bit about some of the sample implementations that are out there. I'll turn over to you, Russ. And think of your questions too please, because we do have this distinguished panel of folks here who will be available to answer questions when we're done with this.

RUSS MUNDY: Thanks Dan. And as Dan said, please feel free to ask questions and if you think of something as we go along raise your hand, don't mind interrupting. We can adjust the amount and how much talking there is in the slides. So ask questions going through this part too if you desire.

One of the things that people often don't think very much about is whether or not DNSSEC is an add-on or if it's just really all integrated tightly into the whole DNS structure itself. And the reason that we started out by explaining DNS is in fact that DNSSEC is a very tightly integrated aspect of DNS.

It was designed so that it uses the same concepts, it uses very similar structure, it adds new structures so you can do the cryptographic verification. But in reality it is additional material that's added into the

DNS and so it is an integral part of the DNS. So the whole object of DNS is to have the content that's held by some owner of a zone get delivered to somebody that's asking questions about that zone. How do I get information X? Do I get the address record? Do I get the mail exchange record?

And so we want to make sure that however you're operating today for your DNS – whatever aspect you do of DNS – that you ought to, as you move into doing DNSSEC, consider doing it in a similar manner to how your DNS operation is run today.

So today, if you're a big operator of a registry where DNS is a critical important function of what you do, you're probably going to have a knowledgeable DNS staff available either organically or under contract or in some manner. Because DNS is an important aspect of your business and so you want to make use of that staff for doing DNSSEC, just like they do DNS for your organization.

If you're an enterprise kind of organization, hp.com is my example here because HP is a very large organization an in reality they've been a big DNS player in many ways for a lot of years. And for them to do something with DNSSEC, they would want to integrate it into their DNS operation in HP. So if you're a small business and you have a registrar that runs your DNS for you, you're probably going to want your registrar to also run your DNSSEC name service for you.

So if you're a DNS user you're going to want to make use of DNSSEC but you don't need to know much more about it than what you know about DNS today. You don't need to know a lot about DNS to use it. So if

you're focused on the DNS heavy information usage, then you should, as an owner of zones, owner of donor information, you should make sure that you give your information the same kind of concern that you do your DNSSEC keys.

You should take care on what's often called the provisioning side; getting your data into the running name servers and making sure that it comes out when it's DNSSEC-signed in the way in which you intended it to come out. So you need to be careful on the side that interfaces, whether it's registrar or an EPP connection – that you get accurate and consistent information.

So you want to have your knowledgeable DNSSEC staff take care that the DNSSEC is done consistently and with the same care that you do of your names and name content itself. So if you're a registry and you have… So starting at the top, the root is of course signed and it's been signed for about three years now, just a little more than three years. .com is signed, .net is signed… In fact I think most of those are signed.

So all of the TLDs you see illustrated there, about one-third of the TLDs there are signed. Enterprise level, down at the next hp.com example you can see within that that that's your zone and that's what the area is if you're an enterprise operator that has a big DNS operation. That's the size of operation you can worry about.

So you've got a lot of internal structure. You should have good DNS staff already that do the DNSSEC things. Verisign obviously is big in the name business and so they're going to have their own enterprise. But if you're your own little business – like I have my own very small consulting

business besides my regular employment; and I run the name servers my self. I am my DNS staff.

So you can have a wide range of people involved. So the two general principles that you want to use is however you're doing your DNS pieces today, whatever those DNS pieces are, look at doing DNSSEC in a consistent manner with that, and you will be successful. There are a ton of tools that you can use but it does take planning.

When you look at the content involved and… Hopefully that's readable. On the far left-hand side it says "registrants". That's actually the entity that owns and registers a zone. The registrar is halfway up there on the left-hand side. And then up at the top is where the registries and the name servers are located.

On the right-hand side is when what we illustrated here in the skit is people asking for DNS information. So the first thing you need to do is know where and what kind of DNS service you have now. Who operates it? How much of a role do you as a person…? Or who's the right person to talk to that's responsible for those functions?

So the functions we're talking about are… You can see the zone data gets generated by somebody – that's the far left-hand side –, it gets put into the authoritative name servers – that was Dan and Roy and myself – and the recursive name server was Jacque and the client was our friend Norm, Joe User.

So it's just laid out the exact opposite. So when they ask the question the information has to first be there. So that's the DNS layout of how it often works. Components can come from a lot of places. The reality is

there's a lot of machinery involved behind the scenes but the user doesn't see it and unless you or your particular direct people are involved in running this big set of DNS layouts, you don't even need to know.

There are a lot of queries, a lot of answer. This is a typical web page filling. It takes about 70 to 150 DNS queries and responses. And any one or all of these can be hijacked, as we showed in the skit. So this is a really big website, which is cnn.com. But remember, it's the zone data that's most important because that's what DNSSEC exists for.

So if you protect your DNS keys more than you protect your zone data you're not giving the proper balance to your security infrastructure that you need to have in place, whether you're doing DNSSEC or not – because it's the content that matters.

And so on the right-hand side you see the greenish… That's an ugly green, it looked better in PowerPoint. Anyway, that's actually where DNSSEC functions, is on the right-hand side. That was the running name server, the validator, asking the question, getting the answer and checking the stars to make sure they match.

So again, do DNSSEC the same way you do DNS and in this case of putting your pieces in, you want to have your signed data, put in by the zone administrator, whoever that may be or whatever function actually does that. So you're including DNSSEC information, just like when we slapped our little stars on the authoritative name servers out here.

And the validating recursive resolvers, he takes his little star – it's a different star but it's used in a way to validate the information. So if

your operation is being ran by yourself, in some ways that's the easiest way to do DNSSEC because you already have the capable staff. But if you're not then you have to look at whoever is providing it. For doing It yourself, if you're using open-source tools there is a large number of open-source tools available that do full DNSSEC implementation now.

The main open-source set of name servers are all DNSSEC aware and there are also – for those operating your own name servers – commercial products that do DNSSEC also, including Microsoft's newest server release; that also does the full DNSSEC implementation. There are other products also. So you can go to your vendor and say: "Okay, I want to make sure I have the latest releases that does DNSSEC properly." That will get you what you need.

Now, if you're using… There are signing services available, so if you don't want to upgrade your software you can do what's called a signing service. There are several of these available in the world and some are free, most are very inexpensive if they do charge you anything at all. And you have an agreement with them where you can ship them your DNS content information, they will do all the DNSSEC stuff for you, send it back and you can load it into your name servers and go.

You still have to have your name servers be able to properly answer DNSSEC queries, but you don't have to worry about doing any of the key management or anything like that. You do hand your zone information to someone else and let them sign it, but a lot of people have that as part of their operation in some ways anyhow.

If you're doing a mixture of these you can have your answer be: "I'll just fit it into my operation. Whatever the right pieces are I can find them easily and put them straight in. In some ways that's easier than if you're outsourcing in some manner of another. If you're outsourcing you're less likely to find a cooperative vendor, unless they're already doing DNSSEC.

But you need to ask for it if they say, "No! We don't do DNSSEC," because that's something that's been a problem for a number of years; people saying, "There's no demand." This is truly changed now, but the more people that ask vendors for it the better. If you're the owner of a name but have somebody else doing all your work, that's the set of things that you're going to want to look at here.

Who is doing your work and who is responsible for interfacing with these people? For the purposes of our simple illustration here, they should do all the work. There is a big mixture of how much is done by an organization and how much is outsourced, but given that they do all the work it's simply to say you run this as a DNS-signed zone.

On some cases some registrars that provide DNSSEC service, it is literally as simple as checking a box on a web page. Others it's not so simple because they may not do it or they may be resistant or reluctant to help you. So that's the essence of how you can go about doing this at a very high level and how it works. So we've got about 15 minutes yet of our scheduled time. Is it? Oh good! So we have lots of time.

So, questions? Yes, please.

AUDIENCE MEMBER: Hi. I'm not an expert on these things and I have an intuitive question. If a middle man can stand between the ISP resolver and bigbank.com, how does the middle man stand between the root and the ISP? The bad guys can pretend to be in the…?

DAN YORK: So you're asking if a bad guy can get between the ISP and Joe User? Is that what you're asking?

AUDIENCE MEMBER: And the root.

DAN YORK: You're asking, when Joe User asks the question of the ISP…

AUDIENCE MEMBER: The ISP asks the root, right? If the middle man gets between the ISP and the root, what will happen?

DAN YORK: We're talking more here about DNS hijacking and getting a man in the middle, but you're right, there are other root hijack where you could hijack the root, which is another aspect of security that this does not protect against. There are other mechanisms that are being looked at to deal with root hijacking.

AUDIENCE MEMBER: Not to the aspect of DNSSEC, right?

DAN YORK: Right, yeah, there are other technologies there.

AUDIENCE MEMBER: That was my first question. My second one is, a new technology, [inaudible 00:57:30], faced a problem of deploying them quickly to take the place of the other one. My question is, does DNSSEC have some incentive to support its incremental deployment?

DAN YORK: So you're asking about the status of deployment? Actually, on that note I'll mention that if any of you are very interested in DNSSEC, on Wednesday over in room 3B we're going to have a… Sorry, 3C, I apologise. We're going to have the DNSSEC deployment workshop, which goes from 8:30 in the morning until 2:45, and we have a whole series of more technical topics around this.

But at the very beginning Steve Crocker will actually be here talking through a series of deployment maps showing the status of where the deployment has gone over the past couple of years and the status of where it is right now. As you'll see in those maps, when he shows them, there is a lot of deployment happening around the area.

To answer a couple of specific things about the status of it, there are a number of large ISPs around the world who have deployed validating resolvers; the largest perhaps being ComCast in North America, who's

turned that on for their 18 million and something customers. A number of European ISPs have turned this on. The biggest one of late has been Google, who have enabled DNSSEC validation for their public DNS service.

So anybody who's using Google's public DNS resolvers now gets the protection of DNSSEC. And we're seeing lots of signing happening as well.

SPEAKER: As Dan said, Steve's maps give a good geographic view. The count percentage wise is still low. The count is in the millions of zones that are signed.

AUDIENCE MEMBER: How much of the deployment of DNSSEC…?

DAN YORK: It's hard to answer that precisely. Most of the gTLDs are all signed and about 100-odd of the ccTLDs are signed right now, including most of the largest ones. So we're at a point where out of the 313-odd domains that exist today, about one-third of those have all been signed at the top level, the TLD level.

Now, below that we don't have a whole lot of information on the second level in all places. We do know for instance that in .nl, about 20% of the domains in .nl have already been signed. In some of the other areas like .com, they're not. .nl is much smaller in that regard. .cc has a high

percentage of domains that have been signed as well. And we keep seeing more of that climbing in the various different TLD throughout there.

AUDIENCE MEMBER: I will pay attention to the workshop that you mentioned. In general, what is the main challenge in front of the deployment of DNSSEC?

DAN YORK: What are the main issues?

AUDIENCE MEMBER: What are the challenges?

DAY YORK: The challenges? The two main challenges are getting the number of resolvers out there deployed that are doing it. One of the things that we've been talking a lot about within the projects that we're all involved with is that it's become very easy to turn on DNSSEC validation, because sometimes it's just one line of code that has to be added to a config file and the validation can start.

So part of what… That's one of the challenges; letting people know about the value that DNSSEC can bring in prevention man in the middle attacks. There's something else we'll talk about on Wednesday called the DANE protocol. And if you were in here before you'll have heard me mention it with regard to SSL and protecting the integrity of TLS

certificates. There is something called DANE that is part of that. And DANE needs DNSSEC.

So we're starting to see… One of the questions is, why deploy? Why do it? And that's one of the things that we're starting to get really good answers with because we're able to say: "Look, you can use this to provide a trust layer to the Internet." You can provide a much higher degree of integrity to the answers. So, why? What's the reason for doing it is one and two is getting the actual validating resolvers out there.

Google turning on their public DNS was a huge step forward because people said: "Look, it's good enough for Google, Google's enabling DNSSEC validation for their public DNS servers. They do it at a huge scale. If I'm an enterprise why not do that?"

And as an ISP why not enable it there, because if people are really looking for security, ideally as an ISP you'd like them to use your DNS servers than potentially go off to Google's public DNS servers. And it's a bit more secure if you really want to look at it coming from the local ISP. And the third part is encouraging signing, because we need to have the two pieces; we need the validating resolvers and the signing.

And you heard earlier today somebody stand up at the mic and ask Google why they hadn't signed their zones. And that's part of the dialogue that we're continuing to have with the major players. A lot of organizations have gone and signed their zones. A lot of companies are doing that. We're seeing an increasing exposure there. But those are the three main ones.

For a while it was getting software that would support it too and applications, but we've seen a large range of libraries now that support DNSSEC, and so the fourth would be getting more application developers to include DNSSEC validation in their apps as well.

JULIE HEDLUND: Excuse me one second. A couple of little process things. Please state your name, always, any of the questioners. Also please make sure you're speaking right up into the mic because people in the Adobe Connect room are really having a hard time hearing the questions. Thanks.

MARK BUCKELL: Hi, my name's [Mark Buckell? 01:03:53] and I was hoping for some more technical explanations about what you were doing, but I was just wondering how much DNSSEC adds to the overhead of the query, time-wise, because I have my main domain that I'd like to have secure but i have sub… Hosts underneath it for DNS black lists that I have that I provide for spam filtering companies, where basically performance is more important than security because I don't believe anybody would ever spoof that.

Can I isolate part of my domain to be secure and another part of the domain to be insecure? And what would the penalty be for having this security performance-wise?

SPEAKER: If your question is about an individual zone that has some delegations that are signed, that are delegated and some delegations that were not signed, was that…?

MARK BUCKELL: Yes, and the total penalty for processing DNSSEC.

SPEAKER: So the technology does support that approach with what's called NSEC 3 and opt-out. Roy, you've probably had more experience with this than anybody. Do you have any performance information?

ROY ARENDS: To be fair you don't need NSEC 3 or opt-out to have some delegations signed and some delegations not signed. For instance in the root zone not all delegations are signed. But it's perfectly possible to have some information there that's not signed. In terms of overheads, if you sign a zone it of course depends on the size of the key and how often you want to do it.

Before you publish, also, you need to sign it. There are some overheads, absolutely, but if you don't change the zone that often they are marginal. If you do change the zone every few seconds because you serve black lists to companies who need them, and it's highly dynamic and changing, etc., it's going to cost a lot to get an infrastructure ready to use sign.

MARK BUCKELL: As I say, so you have to resign it every time you change or…?

ROY ARENDS: Yes.

MARK BUCKELL: Oh, that's not going to work at all.

ROY ARENDS: Yeah, every time that you change the records you need to resign it. But the actual latency of the query, etc., when you're going and sending the request your resolver is sending a bit in there that says: "Give me the DNSSEC components to it too." So it's getting one response back that has the additional records in part of it.

It will increase the size of your DNS responses, because you have additional records in there. But speed wise that's what you're incurring. There are operational issues and if you're doing a lot of constant changing of your DNS zones then you will incur a cost operationally in resigning that every time that you change the zone data.

I would mention to, just as a note, for folks who are in the back or who are leaving, do pick up the sheet if you don't have one, of the session notes. It has a number of URLs on there that would do it and where you can go to to learn more. Mark, to your question about the technical side, this is the DNSSEC for Everybody Beginner's Session.

Wednesday is the deep technical session, and believe me, some of the sessions we'll be doing on Wednesday, we're going to be diving deep.

So if you're looking for technical information, Wednesday's a great day to be there in 3C. Yes?

PAUL MACHENE: My name is [Paul Machene? 01:07:42], I am an ICANN Fellow. My question is regarding DNSSEC and implementing it from the registrar's perspective. So, you have all these clients for whom you've already created the domain names and you want to apply DNSSEC. So I want to know, from your experience, what's the best way to go about this without actually breaking the domains or without actually having operational problems with the keys?

DAN YORK: Sure. There are a couple of answers to the question. We have to think about the registrar portion of what you do, versus the DNS hosting part of actually operating the DNS name servers, because there are two different parts that come into play – between the registry and the registrar there is a passing of information that builds this global chain of trust that's there.

And separately there is the DNS hosting component, which is the part that's doing the signing of the domains and the pieces that are there. So there are pieces to think about on each of those and there are some White Papers out there that we do have, that provide some more detailed information about the parts with each role. Does that help?

PAUL MACHENE: Yes, I understand that.

RUSS MUNDY: Especially from the registrar perspective, as strictly doing the registrar function, all that's really required is to be able to properly handle DNSSEC related information that you pass from whoever is operating the name server from that zone to the registry that you're working with. So from the pure registrar/registry function it's fairly confined and straightforward.

It's being able to process, store, properly handle and pass the information through. It's completely distinct from operating the name servers for that zone or the name servers for the parent. And that's what Dan was getting at. But there are two separate pieces. Folks tend to wrap them into one set of things. But they really are separable even though in most places the same organization tends to do both.

PAUL MACHENE: So it means basically… Assuming you own the name servers… It means you are the one responsible for the technical operations of those name servers in respect to DNSSEC?

RUSS MUNDY: Yes, that's right. So as I was saying earlier, if the registrant, the owner if you will, of a name, can have as simple of an interface as, "Check this checkbox on the web interface," and then the operator of the registrar, who also happens to operate the name servers for that zone, do all the technical pieces back behind – invisible to the user; the user doesn't need to know or see anything else – are all DNSSEC signed. It does

everything internally with one important exception and that is the passing the DNSSEC information upward to the registry above it.

PAUL MACHENE: Okay. And do you have examples of maybe providers who are maybe giving their clients or the registrants an option like that, where they just check…?

DAN YORK: Yes, I'll give you a couple. Before I do that though, let me just say: if you're operating your own name servers, most of the major name servers are out there, whether it's authoritative name servers, whether it's BIND, PowerDNS, Microsoft Windows server, any of those, they all have now made it very simple to go and enable inline signing or automating that signing component very easily so that you can just enable that in the server and it'll do the signing automatically.

There are projects like the Open DNSSEC Project, which is a little bit more involved but it provides additional levels of security that also work with that. There is a series of DNSSEC tools, there are a number of libraries and other pieces that can help with that. So the tools are out there to make that work. As far as the user interface, there are a number of examples.

Now, I should say there are some registrars… I'm trying to look to see if Frederico's here now and he's not… There are some registrars that automatically sign. When users come DNSSEC is just part of what they do, so that the end user doesn't even have to worry about it. They

register a domain with that registrar and DNSSEC is just part of what that registrar offers. So it's just there. [music playing]

Other ones like GoDaddy for instance has a very easy thing. They have an upsell for if you buy their premium DNS service, you wind up… [music playing] [laughter] We've got some kind of dancing or something going on here… If you buy their premium DNS service then what happens if you get a little box that has a checkmark and you just check that and it takes care of the domain.

So it does all the signing for you, it rotates keys, it does everything else. DINE, another company, they run DINE DNS or DINE ACT and their services, they've made it… It's not a checkbox but you go and enter a couple of buttons and you make a couple of choices and hit okay and it does it all for you. There are a number of other user interfaces out there that are similar to that but part of that is, Roy, making it simple.

ROY ARENDS: What?

DAN YORK: Names Beyond is another one that has made it relatively simple. So the examples on there just more people need to be doing that. [music playing] Obviously the next room over there is having a little bit more entertainment than you are here – because it sounds like they've got some kind of dancers going on over there… [laughter] What? It's outside? Oh. Speak loudly over the drums and dancing feet and shrieks.

GUILLEIRNA: I'm [Guilleirna? 01:14:20], an ICANN Fellow from Pakistan. I am a beginner, a basic beginner. So in this skit there were two cases. In the first act the ISP wasn't looking for the signatures but in the second act it was. So do the ISPs always know that they have to look for the signatures?

DAN YORK: The ISP or whoever is doing the resolving does. In our case we showed the ISP as the validating resolver. It could have been your local computer or your mobile device or something like that. I run a validating resolver on my laptop because I'm a DNSSEC geek, but I do that, and so when I go to it my local resolver sends out its queries saying, "I want DNSSEC information there," and it does that.

Most people's laptops in here don't have their own resolver here so what they do is go and send it to the ISP and the ISP then goes off and does that. And if an ISP supports it they're typically going to turn it on all the time. ComCast turns it on all the time on their network. Some of the ISPs in the Czech Republic and Sweden and others have turned it on all the time.

You're probably not going to have situations where you go "partial" or you have ask for it. Although when Google first rolled it out on their public DNS servers they made it optional that you had to request the validation when they were first testing it. And there is a way you can do that. And then you have to use a special application that sends a request and does that.

So you can have a partial implementation like that that's optional, but usually ISPs are going to just say: "From this date we're now enabling DNSSEC validation." And that's it.

GUILLEIRNA: Okay, I have another question…

DAN YORK: By the way, on that I'll also mention there is a great White Paper from Surfnet out in the Netherlands, where they've gone and talked about how to go about enabling validation in name servers and they show examples from three of the top name servers that are typically deployed in name servers. That's a good one to check out.

GUILLEIRNA: Surfnet? Okay, thank you. The other question is, Dr. Evil was able to come forward just because of his speed, right?

DAN YORK: It was speed and the fact that… Speed and he got to the right point in the network where he would be able to go and do that. So you could think of an open Wi-Fi network, like for instance this ICANN network… We should do that as a… So if he got onto the network and he saw the query that was coming from the ISP, that was being sent out there, then he could respond back and do it.

So it's speed and it's getting to the place in the network where he can do it. Oh, and Roy wants to say something.

ROY ARENDS: Thanks Dan. I'll use this one in a minute for a next one. The additional thing is, it's not just speed but it's also plurality. For instance, in a specific Kapinsky attack, what you do is you send a lot of answers, as an attacker, and hope that one will actually make it. So it's not just speed. Yes, you want to be first but also what helps an attacker is to send multiple answers in order to get there first.

GUILLEIRNA: Okay. When you get a validation, is it related to speed and how to… If you're getting a validation, can you try and get the accuracy or something related to that to get to the finder, the ISP in that case, at the right moment? Do you understand?

ROSS MUNDY: I think so. Do you mean, can you just hit it at the right time and get it? Well, the right time is really whatever the first answer that arrives at the place the question is being asked from. So DNS software that is going to get answers a lot is designed for efficiency, and once it gets one answer it says: "I've got my answer, I'm handing it back to whoever asked me and throwing it away."

Now, it could be the technique that Roy just described where a whole bunch of particularly formed queries… The reason you saw caveman Kaminsky in the slides is Dan Kaminsky identified a way to exploit a weakness in DNS that would, at that time, allow you to compromise – in

a manner of usually less than a minute, sometimes a small number of seconds – most of the running resolvers that were used on the Internet.

So many of the things have been changed since then so it's not as easy to do but it still can be done. So however you get that answer to the machine that's asking the question, the first answer, if it's accepted, wins. And the thing when you use DNSSEC is the answer will not be accepted unless it passes DNSSEC validation.

DAN YORK: And note too that, again, this validation includes not only validating the specific answer that comes back, but validating that the answer that comes back lines up with the chain of trust all the way up to the root. So even if somehow, mystically, an attacker was somehow able to get a signature that might look good or something that could come back to there…

They could do that relatively easily, they could sign the record and generate something to come back. So it would look like a valid signature but the resolver would say: "The signature looks good but it doesn't have a chain of trust going all the way up to the root. So even though it looks good, it's bad. Get it out of there."

GUILLEIRNA: That's great. Thank you.

EDUARDO: My name is [Eduardo? 01:20:56], I'm from [Poland? 01:20:57]. I'm a Fellow. My question is related to the application itself. What happened if the resolver's application is vulnerable and they attack it and attack it and our IP [resources it? 01:21:13] the resource records that point to the root? Are they [intrusions? 01:21:16]?

DAN YORK: I'm sorry. You're saying for the application, what happens if the attack…?

EDUARDO: If the application is vulnerable to attack?

DAN YORK: Vulnerable? Okay, sorry. Okay. Yeah, so…

EDUARDO: Is there anything that has been invented to prevent this sort of thing, in terms of the resources themselves, from being [inaudible 01:21:41]?

DAN YORK: So, DNSSEC protects the integrity of the DNS answers and makes sure it gets back there. That's the problem is solves. It does not protect against other types of things, such as an attacker taking over the resolver. If someone were to compromise the ISP's servers, for instance, and turn off validation in the resolver, then when Joe User went to request it the

resolver would go out here and be getting the DNS information back and not asking for DNSSEC results.

So the attacker could still go and do this. Similarly, an attacker could compromise Russ, if he was the bigbank.com name server. If an attacker could get in there, get root on that server and change it around, the attacker could conceivably just remove DNSSEC so that there were no signed responses – although there is a check for that. But he could do other things to, like change the records and regenerate the keys.

So DNSSEC does not solve the… You still need all the other pieces of IT security, but it solves this issue about, am I getting back the same information that was put in the DNS? That's what it's about. And note too, we'll talk on Wednesday about this thing called DANE, where some of that information stored in there could be SSL certificates. And it could be other things. Anything you put in DNS could be there. Yes?

ALEJANDRO ACOSTA: Hello. Alejandro Acost, I'm also a Fellow, from Venezuela. About a year ago I turned on the validation in my DNS server and so far so good. But my question is, if you know the numbers, how much did turning on DNSSEC impact the processing in Google's DNS software, or something like that? Is it significant?

DAN YORK: We don't have Warren Kumari here to bug. Roy, from a utilization perspective, any sense on what it operates or what additional hit there might be to operating and validation resolvers?

ROY ARENDS: The major penalty hit for validating a resolver is the actual validation. It depends per validation on the size of the key, which defines the size of the signature. But there's a great deal of optimization going on. After something is validated as being correct, it's then cached. So it doesn't have to be validated again.

For every cache miss… Sorry, a cache resolver can either have a cache miss or a cache hit. If it's cache hit, it's cached, you don't have to look it up. If it's a cache miss it needs to be resolved. It then needs to be signed and the validator can validate it. And that's the significant… Maybe 200 milliseconds for a 1024-bit RSAKEY.

ALEJANDRO ACOSTA: Okay. My question was more oriented. Suppose you have your DNS server running and it's consuming around 10% of the CPU. Tomorrow you turn on the DNSSEC validation – of course it depends on how many [inaudible 0:25:18] per dollar – but how much [inaudible 01:25:21] do you think it can impact your CPU usage? 10%? Maybe I would need to change my server because it really impacts the CPU usage?

RUSS MUNDY: Funnily enough, this depends on which country you're in. If you're in the Netherlands, .nl is vastly signed. And if your ISP switches on validating from day zero to day one then yes, it has a massive penalty hit. I don't know what that is in numbers. If you are in the UK however, we only have about 1,500 domains so far, out of 10 million domains,

signed. So if it's in the UK then 10% it will stay. 10% [inaudible 01:26:10] switch DNSSEC on.

So yeah, given that local… For instance in Europe, and I believe in many other countries, local ISPs tend to resolve mostly local domains, as in Netherlands – .nl – and UK – .uk – so I hope this answers your question a little bit.

NORM RITCHIE: Well, a couple of years ago JPRS undertook a study, which I never saw the final results from. The preliminary results that they briefed at one of these meetings a few years ago indicated somewhere around 10-15% seldom didn't hit it at 20%, depending upon exactly the factors that Roy was mentioning. So it's a hit.

ALEJANDRO ACOSTA: Thank you so much, and the answer was great. Thank you.

DAN YORK: Yeah, and there is some additional work being undertook in a number of different places around these kinds of questions. So I guess we have a question from the chat room?

JULIE HEDLUND: Thank you very much. Guest #2 in the chat room asks: "Is there an industry standard HSM device recommended for the deployment of the DNSSEC in a low-cost model?"

DAN YORK: Ha! What the person asked was that to do the signing there are different levels of security that you may put around the keys; the actual keys that you use to do the signing. Now, if you're the root you have an extremely secure process that involves using what are called HSMs – hardware security modules – and if you want to see a really secure system the root has this whole process where they use these HSMs, which are small little hardware boxes, inside of other boxes, inside of locked cages and in separate rooms…

And there's this whole big process that people go through to make sure that everything is done exactly like this. And it's videoed, live-streamed, you can go and watch a key ceremony and see the whole process that goes on. For the average enterprise, for the average person who's there, they may choose, depending on the level of security that they require, to use an HSM device which provides this higher level security for generation of a key.

For other enterprises where they don't require that level, you can do software key generation. Some of the other products that are out there, like open-source software projects, will do that for you. So you don't have to have an HSM, but an HSM does provide a higher level of security. There are any number of devices out there that you can buy for the number of zeros that you want to add onto… However much you want to pay.

Now, Rick Lamb was just here in the last session, showing how he built a small little HSM that would do this kind of key generation for a very

small amount of money. I forget what he said but it wasn't much. US $25? Okay. So you could pay $25 and do it Rick Lamb's way or you could pay $25,000 or however much more you wanted to pay for a higher-end box.

So the response is there is no standard. That's the short answer. There is no specific standard. There are many choices out there ranging from various different levels to software-based HSMs as well. Go ahead.

AUDIENCE MEMBER: I'm [indaudible Robert? 01:29:49], I'm an ICANN Fellow from Uganda. My question is, in the skit you are showing where you have a complete DNS lockup. How about if there [inaudible 01:30:06] in each cache? If you don't have a whole lockup?

DAN YORK: Right. If the ISP had already cached, if it already knew how to get to .com…? Yes. The process still works because what would happen is in that case is… So what we're saying is that DNS is optimized so that Jacques, as the ISP resolver, didn't have to go to the root every time. My role could have stepped out of there because I would give Jacques an answer that he would then cache for a certain amount of time and he would go back to .com and ask the .com name server where bigbank.com is.

And this would be without involving me at all. I would just step out of the picture as the root server. But .com is signed, okay? And it was signed in root, we've all agreed with this. So Jacques would have gone

here, got the answer from bigbank.com and then it would have tried to do this validation in the chain of trust.

If the chain of trust was broken then Jacques would have started out again back at the root, going back and resolving it. Because somebody could pretend to be the .com name server. Dr. Evil could have injected that record and said: "This is the .com record really," and gave a signature to it.

So there is a possibility in there but this global chain of trust protects the caching so that even if Jacques knew bigbank.com, if it had that cached, if he went there and asked bigbank.com, bigbank.com gave him an answer for www.bigbank.com, then Jacques would check the chain of trust and say, "Does this work?" And he's already cached the other records for the root and .com and he could check that and see if it works.

If it doesn't he could say, "Oops, there's a problem here," and perhaps restart the process again to do it. So the caching totally works and it lines up with the chain of trust. Other questions? This has been great. We've had some excellent questions here today. One more?

MARK BUCKELL: I was wondering whether there was any easy way to just pick up the signature information without actually doing the DNS query? The reason I ask that is I thought maybe it's a spam filtering technique that domains that had signatures would probably be more trustworthy than a spammer that just grabbed a domain yesterday and probably wouldn't

set up signing as a way of putting false positives, you know? Do you have any thoughts on that?

JULIE HEDLUND: Excuse me, could you introduce yourself please sir?

MARK BUCKELL: My name is Mark Buckell, junkemailfilter.com.

SPEAKER: Thank you Julie. Yes, and there is, and it's easy. You can query for DNSSEC records and get DNSSEC records if you want them. When you do the DIG command you can just say "DIG + DNSSEC" and you'll get back the… And you can query on the DNSKEY record for instance, which would be the key that's in there. So you can do that, or you can query on RRSIG, which is another one that's there. But that's all.

So you can use it with DIG or any other library that has DNSSEC support.

DAN YORK: All right. Any other questions? Okay, well I want to just wrap this up then. I'll say thank you very much for your participation here. Again, this is the workshop that we have, the DNSSEC for everybody. Again, Wednesday morning, 8:30, in room 3C we'll be having the DNSSEC workshop, which is a whole series of sessions. If you're curious about deployment I'd encourage you to come at the beginning because Steve Crocker does have some great maps and we'll talk a bit about the overall view of it.

Then we have a panel on DNSSEC in Africa where we'll have a number of folks; I see Mark Elkins is here and we'll see a number of different of people who are involved with DNSSEC in Africa, we'll talk about what's going on with DNSSEC deployment in this region. We also have a presentation from Patrick Fälstrom, who will be talking about the roles of the registrar, the registry, the hosting providers and the various different players in that and how the whole system works.

We'll also be talking a bit about… Russ and I will be back talking about some of the innovations with DANE and some of the ways in which we can use the higher level security and applications, and a range of other topics will be part of that. So we encourage you to come. Again, pick up that sheet that has the URLs for more information. There is the DNSSEC tools project, the Deploy 360 program and the DNSSEC deployment project.

And I want to thank as well Roy Arends, Norm Ritchie, Jacque Latour and our Dr. Evil has disappeared on us; he's gone. So if you can all join me in a warm round of applause for everyone? Thank you very much. [applause]

[END OF TRANSCRIPT]