Julie Hedlund: Welcome everyone to the DNSSEC workshop. We're going to get started in just about a minute or so. Welcome everyone to the DNSSEC workshop and thank you so much for joining us here today. We are going to try to proceed on time so we'll be following the program pretty closely. And if you are on Adobe Connect for any reason, where we are running the program simultaneously, I just would ask that you mute your computer otherwise we will get an echo from the sound coming through on Adobe Connect into this room as well.

I'm Julie Hedlund and I want to welcome you here and I'm going to go ahead and turn the program over to Steve Crocker who is the co-chair of the DNSSEC Deployment Initiative.

Steve Crocker: Thank you Julie, thank you very much. Welcome everybody. It's a real pleasure to be here. We have been running these workshops for quite a long time and one piece of homework that I didn't do but I will do next time is try to count up how many we have been doing because I think they've been running for five, six years – something like that.

Traditionally in all of the time that I've been involved in security, security has been usually done in dark places, it's usually been last. I've been to meetings in the middle of the winter in Stockholm. And it is truly a delight to be up here in the 32nd floor with windows all around in daylight, well, what constitutes daylight for San Francisco anyway – and

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

in spring time. So at the risk of very bad jokes very early in the morning it's clear that security is emerging out of the darkness.

And indeed DNSSEC is now in its infancy and it's been a long, long haul. So it's a good time in this particular area. These workshops have, require a lot of work so I want to acknowledge the other people who have been central in making everything go – Julie Hedlund on the ICANN staff has been in the middle of all of this and put all the pieces together. Russ Mundy to my right – Simon, is Simon here – there he is hiding back there. Simon McCalla from Nominet; Marcus Travalier from SIDN has been a very active member just recently gone off to do some other things but much of what you'll see toady is the product of a full team of us.

And it's also appropriate to acknowledge our sponsors and you'll thank them particularly when you get your free lunch – free to you not free to them – at lunchtime. So the PIR people who run .org; VeriSign; the Dutch registry SIDN; the Swedish registry .se; the open DNSSEC consortium; GoDaddy; Afilias – I think we got them all – no Nominet – and Nominet of course. It's been very, very helpful to have everybody pony up and make life easy for the rest of us.

A couple of other key announcements – the ccNSO/DNSEC panel discussion today will take place in California West, which is in some other county – I mean, the distances around here are incredible. This session will go until 1:45; theirs starts at 14:00; we have clocked it and it is barely possible to get from here to there in that period of time. The coordinating panel discussion following this workshop – the goal of the panel is to identify if business and marketing factors are relevant to increased registrar and registrant buy ins for DNSSEC; and if so, which of these factors and how can it be influenced by a ccTLD registry. Alyssa Kelly the Head of the Nominet, is the chair of the session and Chris Disspain from the Australian registry will moderate.

We also have two other interesting and relevant presentations that are not included for direct presentation in this workshop but which will be posted on the meeting schedule site and listed under additional resources – one is DNSSEC Validation Measurement – how to count validators by Kazunori Fujiwara from Jprs. I think he presented this in the DNSSEC OR session on Sunday if I recall. And Lutz Donnerhacke- How IPv6 and DNSSEC changed intranets – Lutz is sitting over there. Did I get this right, the title? Yeah? Good. So with that, let me move forward.

Here is the agenda for today – I will stop talking relatively soon and move on to the panel of the people who are up here – Russ will take over and moderate at that point. We have Application Security with DNSSEC and DOSETA. Some distant relative of mine named Dave Crocker will be presenting. We'll have a break and then Innovative Uses as a Result of DNSSEC – this is where we get to look a little beyond the direct implementation of DNSSEC as well. And then Activities from

Around the Region and then after lunch we will have a panel discussion on DNSSEC Signing Services, which are popping up all over the place. The order that's listed there will be changed just slightly – Matt Larson has multiple obligations and will start off the session and then we'll proceed in the order shown. Thank you.

I've been trying to keep track of DNSSEC deployment at the top level domain level, and principally with large focus on the cc level, and developing a set of maps. It's devilishly hard to get the counts accurate. Let me characterize that as good news – that is there is more stuff going on then we can keep track of. So let me show you what we have – you will hear some other numbers in different places and the numbers don't all line up – we're going to work to try to align things, but there's a reason why there are some discrepancies.

I've been taking surveys for a period of time and asking four questions with regard to the deployment at the signing level and particularly at the top level domain level. The last question in the sequence is "when are you fully operational – the zone is signed and you're taking registrations and everything is up and running in a stable level". These are shown in blue on these maps. And this shows the status as we understood it at the end of 2010 – on December 31st 2010. And this shows 27 countries that are fully operational as of the end of last year.

The status just below that in green is partial operation and we show another 15 – these are ones where typically the zone is signed or there is a degree of testing going on that is visible and actively engaged with the public, but not full operation. Typically it means that the delegations are not being accepted from the registrants. But it could mean other things. And where I was eluding to a discrepancy is what we have not tried to do here, but which I will seek to do in the future, is match up these numbers with the numbers that show up in the root.

And the specific issue is this – we count a zone as signed and in partial operation if the zone is signed, but they do not have to have forwarded their key up to the, or their DS record up to the root. When you look at things from the root perspective, there's a very precise count of how many things are in the root at any given instance, so there could be a discrepancy in there. As I say, we'll seek to line that up.

A qualitatively different sort of measure is when has a registry announced that it is committed to getting its zone signed sometime in the future. We've tracked DNSSEC deployment over a long period of time and there's a lot of things that go on behind the scenes. And very often they go on for a long time. There's analysis and study and test implementations and a lot of debate very often as to before the registry decides that it will actually go forward. So the announcement process is significant and it's followed in every case that we've observed with actual deployment after that, but no certainty before the announcement that that's actually going to happen.

So that's show in you have your choice of how you want to call that – orange or light brown or something. And here as of the end of December we showed only two in that status. And then experimental is where there is visible experiments going on but no commitment yet that it will actually move forward into full scale deployment. And we show six there. So this map that you're looking at is, as I said the end of December last year – next slide is end of March this year.

And the numbers are slightly different – flip back and forth if you would for me and sort of animate this. 27 to 30 in that three month period – operational. Let's see what are the 15, 12-15 and five experimental. Now, move forward to the next slide and we get to June of this year. So this is slightly in advance of where we are and the numbers move up – 34 operational; the number of one's in partial operational would be converted.

And now move forward – this shows the end of the year and the numbers are basically the same. I think maybe identical. Then – I think we have one more – an estimate for the end of 2012 – and we're not showing any differences. This is almost certainly not the way things are going to unfold, this is just based on the data that we have. So we try to make this data available and we're gradually working on how to improve the presentation and improve the data that we have.

Any questions with respect to any of that? This will bring to the end my overly long, even though short introduction and then we'll move onto more substantive things. Thank you. Russ?

Russ Mundy: Thank you Steve. And is this – I'll move it just a little closer – there I think I'm being picked up. Well, this is the panel for DNSSEC deployment plans and how they affect products and services. I have not only a strong professional interst in this since my group is the provider of DNSSEC tools and we've had a Firefox fully enabled for DNSSEC capability for some time just recently got it running on Windows, so I'm very interested in hearing some of the folks with the real products are actually doing. We have also DNSSEC running on a number of other platforms like the Nokia 900 cell phone, which I'd be happy to show folks if they're interested at some point.

But because of the short length of time, we are asking all of our presenters to keep it to 10 minutes, and everybody is excited about what they do as I couldn't resist plugging our things that we've done. So what we've done is we've ordered the presentations by alphabetic order of the first persons last name, in the case of two people from one organization. So we'll start right off from the folks from Mozilla – and it's Lucas Adamski and Brian Smith. Please go ahead.

Brian Smith: Okay, today in Mozilla's platform we don't have any explicit support for DNSSEC built into the platform. What we do is if there is any DNSSEC support in the operating system, then we'll use that support if the operating system is configured to support some kind of DNSSEC policy. In particular on Windows, you can have a very complicated policy regarding DNSSEC where some domains are enforced, require validation and other domains don't.

There are in addition as Russ mentioned, SPARTA has done a lot of work to enable Firefox and Thunderbird to require DNSSEC validation using a variety of mechanisms including a simple shim library that basically you load in front of the standard DNS resolver library that provides basically instant DNSSEC support for our applications. And they also have more involved support that has not been integrated into our platform, the trunk of our platform yet.

In addition to that, there are some third party add-ons which I will show on the next slide. One is DNSSEC Validator and the other one is the Extended DNSSEC Validator. So, one thing that I want to note on these is the fact that they have on the Firefox user interface in particular, in the top image you can see that we have a very positive indicator of security with the green bar and then next to it we have a very negative indicator of security with the DNS validation error indicator from the add-on.

And I think that this is a problem that we're going to see a lot in the future that if we have indicators of DNSSEC validation or lack of validation in the user interface, that we're going to be sending these conflicting messages, especially in the short term, and users are going to ignore, start ignoring one or both of these indicators. So I think that that is one of the primary concerns.

And in the Extended DNSSEC Validator case, you can see that in the primary user interface they don't have any indicator. And I think one problem with that is most users are not going to click on our security indicator – the blue box – to see the detailed information about the DNSSEC validation. So, most users are not going to really experience anything different with this type of user interface. So I think that because we are the portal from which the user experiences DNSSEC validation or experiences the internet…

I have my own copy and it looks great so.

Julie Hedlund: They actually look fine in Adobe Connect. Sorry, shall I go ahead?

Brian Smith: Yeah, let's go ahead. My slides are not that pretty. If you can't see them I'll…

Julie Hedlund: Okay, that looks better.

Male: I think the problem is that you forgot to sign the slides.

Julie Hedlund: That's not it. I apologize for this but I can't advance the slides for some reason; it's not working. So what I suggest is that, Brian why don't you continue with your presentation and we'll work this out. Thanks so much.

Brian Smith: Okay, so besides the inherent advantages of validating information that we receive from DNS, there are some other potential benefits to Mozilla's users that we're very interested in, including the efforts that we will be talking about later, or others will be talking about later in the day – DANE & CAA – to work towards preventing certificate misissuance in particular, we have so many roots in our CA program already and that program is only going to expand, and one of the problems with trusting so many organizations is that you increase the attack surface for CA misissuance. And that's what these extensions are trying to help with. Because I don't see, in particular, in Mozilla's CA program us becoming extremely more strict and cutting out a bunch of CAs.

And besides that there are other advantages including potential performance advantages, where if we know from DNS what certificate a website is planning to use we can start validating that certificate and start even doing the encryption process of encrypting the key that we're going to share with the website before we've even connected to that website. And also we have a feature called strict transport security which right now we are delivering through a header through an https connection that tells us that like basically the website is https only.

The problem is that if the user goes to the http version of the website, it has to redirect to the https version. So besides being slower there's also a potential for that redirect to not happen of course and then the security advantages of HSTS are limited. By putting that information in DNS and securing it with DNSSEC, we can not only avoid that delay that comes from redirecting to the http to https, but we can directly go to the https connection. And we can do this not only for http, but for any protocol.

And in Firefox and other web browsers we have a mechanism where we attempt to connect to a website using TLS 1.0 for example, and if that connection fails for some reason then we're going to try with SSL 3.0, even though this is explicitly something that the TLS specification says that we should not, that you're not supposed to do. And the TLS specification has a mechanism to prevent this. There's so many websites that don't implement this mechanism correctly that we have to do this unsafe fallback. And again, using DNSSEC, we have the potential to prevent this unsafe fallback from happening in a secure way.

In addition, in Thunderbird and other email clients, there's a lot of information that's already being distributed through DNS including SPF where we will be able to authenticate the source of email or authenticate configuration information to make it simpler for users to configure their email client, which is not a security benefit but a usability benefit. So basically there's all kinds of benefits – performance, security and ease of use that are potentially enabled by DNSSEC in our products.

Some of the challenges that we face though include basically because we're the portal that the user sees the internet through, any problems with DNSSEC are going to get blamed on our product because our product is the one that's not working. Especially if our competitors do not support DNSSEC and we do, then the problem becomes whenever DNSSEC is working correctly and preventing the user from going to a website than what the user is going to see is it doesn't work in Firefox, but it does work in some other products. And that makes Firefox look like it doesn't work.

Until there is confidence, until we can see demonstration that people are deploying DNSSEC correctly, that people are not letting their keys expire, that people's routers are not mangling all our DNSSEC, mangling the DNSSEC responses or preventing our requests from being interpreted correctly – until we can find a path where we can be sure that a vast majority of our users are able to use DNSSEC, and/or we're able to try and use DNSSEC and then fallback to a non secure configuration in a performant way that doesn't just completely make DNSSEC useless, then it's going to be difficult for us to turn DNSSEC on by default for our users.

And the other thing is, when there is a problem, what do we tell the user? I think especially when you combine DNSSEC with TLS you might have like the case that we see up here where TLS is telling the users that the site is secure and the DNSSEC is telling the user that the site is not secure or vice versa. Especially, I'm not sure how much value the user or – I don't see how much value DNSSEC provides the user in the case where there's no secure channel being created, but the DNS information is being secured with DNSSEC. And for that reason I think that we are not going to be able to tell users for a long time anything about what is happening at the DNS level regarding security or lack of security. At least as far as the base product goes.

As these extensions show and as SPARTAs work shows there's a lot that you can do in Firefox and Thunderbird right now to get DNSSEC support into our products. And we really want to encourage people to build add-ons, build extensions into Firefox that demonstrate what Firefox should do. And if there's some way that we can facilitate this effort and make things easier for add-on developers to create these add-ons or make their add-ons work better, we're very interested in facilitating that work in our core platform. And I think that the best prototypes that people create are going to be the basis for the design of our future DNSSEC support in Firefox when it comes.

Russ Mundy: Thanks Brian and Lucas, appreciate it. And again, apologize for the slide absence here – or stuckage. I think Patrik Faltstrom from CISCO is next in our list here.

Patrik Faltstrom: Thank you very much. Let's see if I can get my own slides.

Russ Mundy: I believe work is underway to get the slides working again on the main display computer.

Patrik Faltstrom: Yep. So, thank you very much for inviting me to talk about DNSSEC. I'm chair of SSAC as you might know, but I'm also employed by CISCO and I would like to talk a little bit about what we are doing here. Second slide – there was a couple of questions asking whether you your operating system takes advantage of DNSSSEC. And that was of course a question that is a little bit difficult for us, as a company that moves packets mostly, to answer. So what I would like to – and you see under all the other questions there – the question is really – I decided to question the questions as you will see. Because I would like to talk about what we are really spending time on and that is to make sure that DNSSEC works for all of you that have DNSSEC validating applications.

Next slide; third slide – so the first question I would like to emphasize the answer to is of course whether DNSSEC is important for CISCO. Let me just say that it is – it is absolutely necessary for the internet. It is important that it works –

robustness, resilience and predictability for any kind of application, innovation on the internet or something – that of course is important for us. On the other hand, there are some changes that are needed in the various networks that we already have deployed in on the net just like IPv6 now, the changes. So DNSSEC, as you will see in a couple of slides, implies that people running networks actually have to watch carefully on what they are doing.

Next slide please. Number four – so, one of the main reasons why DNSSEC is important for us is of course just because many of our products, if not all, are carrying DNS traffic. It's also the case that many of our products look up DNS records. Almost ever router and product you see that we make at CISCO do look up DNS records now and then. And of course, we do take stability and security very seriously because that is one of the things that are important for our customers. No one would like to buy a product that doesn't work or reboot in the middle of the night. People should be able to buy CISCO products, configure them and they should then just work. Fine lines is not sort of good enough for us.

Next slide; slide number five – the most challenging thing with DNSSEC, and specifically what I've been working with quite a lot, like the last 10 years at CISCO with the DNSSEC, has to do with the fact that when you start to run DNSSEC and you encounter problems, and specifically problems that might be connected to the network infrastructure, it's quite often a combination of three different problems. You have a bad network design.

You have a misconfiguration of the boxes that you have in the network. And of course, there is software in the boxed and because of that there are bugs in the software. And for quite a large number of people it's extremely hard to know which one of these problems there is. And I have an example that I will go through in a minute that I am currently working on that actually I've taken so far two months to be able to figure out what's actually going on.

It is quite hard to detect the problem just because what is really happening, as all of you know that is dealing with DNSSEC, there are specifically two things that are happening. The first one is that we start to use even more EDNS0 then before, which means that we have a header in the DNS packet that looks a little bit different from what it has been looking like – it's a little bit different from what it has been the last 20 years and then suddenly here and there people have very old boxes that they didn't even remember that they actually had because they are very often transparent; transparent proxy server kinds.

And they have been looking at the DNS traffic happily for 20 years and then suddenly after sleeping, there is a packet that it's supposed to detect, and it does. So people are surprised and one of the problems that many people have is actually to know that they have a box there and to physically find it. And the next problem when they find the box is that no one remembers the password to it.

I get, inside CISCO we have – it is not uncommon. It happens every week, believe it or not, that we get notifications from customers that have not rebooted their CISCO boxes for 10 years or more. I think that is something that we should not applaud us at CISCO, we should applaud the power companies to be able to provide that stable electricity feed to the boxes.

So anyways, the main problem with DNSSEC and these old boxes that do get various different kinds of things, is that people detect it often by a delay in DNS lookups. And that is something that is actually easy to detect when you actually do DNS lookups and you look with your eyes on the timing or you have some kind of monitoring system, but it’s difficult for your customers or someone that just clicks on a website to know whether it’s actually a DNS problem or not. So, one of the problems that we have at CISCO when we talk to customers is that people debugging don’t know how to debug. I think we have fairly okay standard on people in our tech, but the problem is that people need to, when they have these problems, convince to actually open a case with tech and sometimes in that discussion the delay is sort of, it’s a little bit unclear that it’s DNS. So what I see is that sometimes it’s a little bit problematic and there is an extra delay before you actually end up with the tech people that can deal with DNS issues.

Next slide please. Number six – so one problem that I encountered in February, just via pure luck, and this is unfortunately how it works, is that there was some customers to an IP or a cable TV company in Sweden that complained to the provider that when they went to some web pages, the browser said "I'm sorry I cannot go to the webpage" – they get some kind of like broken webpage thing.

They did reload and the webpage came like that. And of course when you have customers that do those kinds of things and they go to bazillions of other web pages and they don't have any problem, and sometimes when they go to some webpage it works. So it actually took about a month before the customer care in this IP or cable TV company reacted and said this probably actually is something. So it was actually some customer of that ISP that contacted me and said "I tried to call this IP or cable company, they have a broken customer care system that they don't trust me"

So anyways, no errors were logged anywhere in the network; they could not repeat the problem. So what was it? By pure luck, we managed to find the domain name which have this problem now and then. It was actually the webpage for the Swedish Government. And one of the issues there is that I happen to sort of know the IP people that run the website for the Swedish Government so we started to understand what the problem is.

And we have this combined issue in Sweden that we are deploying IPv6 and DNSSEC at the same time. And a government's webpage can be sort of overloaded now and then and they have an incredibly difficult load balancing infrastructure in front of them. So lots of debugging happened on the government side, but at the end of the day we saw well wait a second, it's not on the website end, and it's actually close to the edge where the customer is.

So what we saw was that responses for queries for that specific domain name were so large, so response came back in a fragmented UDP packet that did not reach the full service resolver. So that's what we thought it was, but it's a little bit difficult to do packet sniffing on the outside of a load balancing network on a IP or Cable TV network with sort of millions of customers. And you only want to capture the packet for non cached versions of DNS packets for domains that do run DNSSEC and not IPv6. Trust me, I started to learn how the regular expressions with TCP dump works.

So, what kind of problems did we have? Well, this was a load balancer on the outside of the full service resolver of an IP or cable TV company and the load balancer had an outgoing flow where the response, which was incoming to the load balancer, was a fragmented UPD packet where the fragment came in reverse order. Okay.

The problem here was that we had a fragmented UPD packet using EDNS0 – so the problem is that we have still not, we are running this in the lab, we do know that the box is misconfigured because it's not configured to take care of UPD fragments correctly. We do know that that is probably not how you set up your network and you probably should not have your outgoing sort of full service resolver behind the load balancer.

I have no idea – that guy that designed that might have smoked something. You should do it differently. And the last thing is that there might be a bug in the software, but we have not been able to reproduce this in the lab. So the guys writing the code don't understand what's going on, but the ticket is not closed yet after like more than a month.

So anyways, what I see I think is going to happen here on slide number seven – is that we will have both misconfiguration and bugs and sort of problems in building the network and one very specific problem we will have with content delivery networks. Because in content delivery networks in some other cases, many of the DNS responses are synthesized and what we have to do for content delivery networks is we probably have to come up with a net mechanism of either increase the amount of http redirects in the CDNs or we need to pre-synthesize and sign the alternative responses given back by CDNs. And both of those things are sort of ideas and strategies that we are working on quite heavily;

including the work with hierarchal CDN networks for video distribution that is sort of going on in the IETF.

One thing that all of us, I think, have been worried about is that validation in the client itself will probably be pretty difficult to deploy. And my confirmation is that yes, today the way we built our network, it's absolutely not possible in reality, unfortunately. We have to make sure that first of all we can do validation in our full service resolvers correctly. And people need to carefully look at their networks on how it's built.

Last slide; number eight – so what are we doing? We're working hard on educating customers how to configure CISCO routers and boxes, specifically of EDNS0. I hear still people talking about the PICs as an example of a box that cannot handle things correctly. Let me tell everyone that the PICs code for EDNS0, when I started in CISCO in 2000 that was the first thing I took care of. So all code in the PICs since 2001or 2002 can handle EDNS0 correctly. Okay?

It's an old rumor but I do hear that the horse is not dead. But just because we at CISCO know that that specific product actually is fixed, we probably have some others that are not, it doesn't matter if you – you can say the PICs doesn't handle this because we know we're right anyway; in that case.

So, EDNS0; fragmented IP packets; and content delivery networks – the community needs to learn more how to handle those things and how to configure their networks. We are working hard on fixing the bugs for going into these things and as I said hierarchal CDNs. And now of course all of you ask so what are you doing on the resolver and DNS service side regarding DNSSEC and unfortunately it's the case that I cannot disclose what we're doing, but let me just say that validation of DNSSEC signed responses is a very high priority on the list of new features that we would like to see in many products in the world. Thank you.

Russ Mundy: Thank you Patrik. And again, apologize for the lack of slides, but I think Julie has gotten another machine – oh that's not Julie. Anyway, someone else has got another machine working. So hopefully we will have slides. And David Lawrence is next.

David Lawrence: It looks like they're close but I know we're short on time, so I'm going to jump right in. I'm a Principle Software Engineer for Akamai Technologies. And for those of you who don't know Akamai we're a distributing computer company. We originally started as a content distribution network to provide increased website performance by delivering the content to users from the servers near them, but in a little over a decade now we've grown to more than a dozen different services that include everything from simple DNS hosting to complex traffic management to a wide area application acceleration. And all of these services depend pretty heavily on the DNS. And we actually are running on more than 84,000 servers around the world in 72 different countries.

So, for our authoritative name service we use an internally written program called Swans, which stands for The Swiss Army Nameserver. It runs on tens of thousands of our servers around the world and it has several different modes for determining DNS answers. Each of the servers is tailored to handle the demands of the service that it's supporting and the simplest mode serves traditional zones much like binds or unbound's and so on. And the more complex modes potentially have millions of updates per minute – far, far more than a traditional zone would. So the most efficient method for returning signed authoritative answers is going to vary by the mode that's providing those answers.

Our enhanced DNS service is pretty unfortunately named because most people then refer to it as EDNS and those of you who are familiar with the DNS know it means something quite different to us protocol engineers. So there is some confusion over that – I apologize on behalf of engineers; we didn't have any part in naming our EDNS service. It is our first and currently only Swans mode that supports DNSSEC.

It provides full service signing of zones complete with Akamai doing all the key management that's necessary. It will also just serve the zones that are assigned by the customer as they get transferred to us. This was implemented due to a mandate from the US Federal Government that all the .gov zones implement DNSSEC. Unfortunately there has been a very low rate of adoption by government customers, we only have a few who are currently using it, because there's a perception among many of them that there's really no consequence, negative consequence to them to having ignored the mandate whose deadline has come and gone to have DNSSEC on their zones.

So, seeing this low uptake rate actually provides very little incentive to us to cover other modes. There is something in the mandate that says – originally it was the gov zones that were supposed to be signed and then it was supposed to progress that any of your services that are provided, even via non gov domains were also supposed to be secured.

And since we host many of those government services presumable we should start signing some of our other zones, but there's been actually very little movement from the government to see that happen. So, one of the biggest lessons that we see here is if the community wants to see DNSSEC advance, it's not just a matter of convincing Akamai that hey this is a good thing; it's a matter of convincing the customers who would be using it that they should be demanding it.

So, in EDNS one of the problems that we also saw was that even to be a full server signer where we do all the key management, it still requires more regular customer involvement then would have been true under a traditional non-sign zone. That’s because secure delegations need records updated in the parent zone and under the traditional registry/registrar/registrant model, there was no explicit operator role. I was an early participant in many of the DNSSEC tech workshops where we just kind of always assumed that the operator of the DNSSEC zone was either going to be the registrant or perhaps the registrar, which because they were providing DNS services. And I don’t think anybody really conceived of the possibility that you might want to have a third party operator of the zone taking care of these things.

So with non DNSSEC having Akamai host your zone was largely acquire and forget. You updated your NS records and aside from managing your own zone, you didn’t really have to do any additional involvement with your registrar. But now with DNSSEC, the customer needs to be more involved because anytime we roll the key signing key you have to interact with your registrar to get those new records published in the parent zone.

So Akamai of course doesn’t have any formal relationship on behalf of the customer for their registrar. And having the community recognize that an independent role would make things far, far easier than having us try to now establish business relationships independently with every registrar; you know, if there were a standard process to say okay we recognize that you’re going to have somebody else providing these services for you so we wouldn’t have to negotiate that every time.

So our content distribution services – this network of tens of thousands of servers around the world – you can imagine there is a lot of records involved in that. So it's very highly dynamic and that makes it really difficult to sing. Our answers are generated from many different variables.

We take everything from network latency to network policies, individual business relationships, network costs – and so with all this high churn and the different answers that could possibly come out on the fly signing become computationally very, very expensive to worry about giving a different answer to each of the different clients that are contacting us every other second. So one of the other possibilities would be to pre-sign all these records, much like you would do with a traditional zone.

Unfortunately, with our dataset, with the numbers of permeations of possible answers we could give there are two factors that come into play. If you were to sign them all you end up signing far more records then you really need to sign – unfortunately you don't know in advance which ones of those you're actually going to need – and with them all signed it actually becomes storage prohibitive. It's far more records then we can keep in memory and if you look at some of the services, the number of records become far more than you could realistically even keep on disc.

So, you can imagine where we'd do some hybrid of on the fly signing with pre-signed records for the ones that we are pretty sure we're going to need to give the answers to. But that's also very algorithmically complex to do so that would take its own special investment in engineering time to solve that.

And all of the solutions mean that we're going to need additional hardware investment to maintain resiliency because they all make the CPU cost and the storage cost much greater. So if we want to provide the same level of defense against network attacks and so on, we'd need an additional capital investment in the servers and memory and so on to support that.

For our other services that we provide – the application acceleration and so on – there's varying degrees of complexity there. It does depend on the mode. None are quite as simple as being able to pre-sign zones the way EDNS does, but it's probably also not quite as complex as the content distribution services. But nonetheless, figuring out the best way to handle that will also require significant engineering effort.

So at Akamai we see significant challenges to advancing DNSSEC – for one, the lack of that universally recognized DNS operator role hampers deployment because there are a number of customers that still would be quite happy just to have us do all of their key management for them but they don't want the additional hassle now of working into their own operation chain how they can make sure that they're keeping their parent records updated properly.

The efficient signing of highly dynamic zones is a financial barrier both in terms of the investment in engineering time and the additional capital outlay for the computers to support it. And most notably, customers still don't really see much value in DNS spoof protection. So, some of them have indicated that even if they did have a signed zone, there are so many resolvers out there that are no validating that there really isn't that much value in signing itself.

So if you ask customers, and I think many of us have this experience, whether they'd like a DNSSEC; they're quite happy to say oh more security for free – sure. But they're less interested in really connecting that value proposition to okay, well you might have to spend a little bit more to really see this happen because we have a number of engineering priorities and if you're not actually caring that much about it it's hard to really get that on our priority list.

And so that's essentially the view from our end.

Russ Mundy: Thank you David. So next we have Patrick Naubert from Xelerance.

Patrick Naubert: Yes. Again, my name is Patrick Naubert; I'm the Chief Technology Officer at Xelerance. Thank you again for inviting us. Normally this presentation would be given by Paul; unfortunately he seems to be busy at this moment, so I've been assigned to it. And since I'm way too chatty, and it's been made clear that my presentation must be below 10 minutes, under pain of death, I'll be reading the presentation just to make sure.

So Xelerance is a DNSSEC product vendor and I'm here to talk to you about our DNSSEC related products. So we're on slide number one – actually we're on slide number two. So, our product family DNSX includes Signer and DNSX Resolver. These products are based on proven open source technologies, to which we contribute back extensively. Slide two please. But first a quick word about our company. We're based in Canada but our market is international.

We are deployed in basically every type of organizations – TLDs, universities, government, private sector. We do a lot of active technical development in a long list of contributions such as LDNS, NSD, Unbound, SSHFP, Perlsnet DNSSEC, [OTO Trust, the NS Perf], a few crypto modules as well in Perl. We're also co-maintainer for Bind on Fedora. We participate in the IETF, DNS-OARC, also The Department of Homeland Security's DNSSEC Coalition Initiative.

Our team boasts a very large technical DNSSEC knowledge pool. And we've been closely involved with RFC 4641BIS. We've presented at many venues and we've won best Security Appliance Award at [FOSI] in Washington, DC last year.

Slide three – our signature product in the DNS world, if you'll forgive the pun, is DNSX Signer. This appliance manages all aspects of your corporations DNSSEC cryptographic operations and procedures. The focus is on simplicity. You interact with the appliance either via browser based user interface or we have an extensive API layer. You can easily hook into this AI with your own product or operational scripts. We even have an iPhone app as a demonstration of this API. Now, as we saw in a number of cases in the near past, when some TLDs incorrect signing procedures broke their zones, active monitoring is essential. That function is a unique part of our product.

Also, our product functions as a bump in the wire. So we integrate seamlessly with all DNS servers including Bind, NSD, Power DNS, and Windows DNS. So introducing Signer in your infrastructure is pretty simple. A new zone can be added in less than 60 seconds. In fact, if you have a text list of all your zones that you want to add that can also happen in less than 60 seconds – we'll see an example of that in a later slide. Slide four please.

So in more detail now – DNSX Signer is proactive in the way that it manages parent relationships and child relationships. We check the zones beneath us and above us to make sure they're right. We have full automated key rollover. We set the key rollover interval to the industries best practices or you can ignore it altogether and just do everything manually. We can generate and submit your delegation signor record directly to your registrar and track the DS record to your registry.

Because we track everything we can never do a rollover in a way that will trash your zone. We only do rollovers when the DS records at the registrar and the registry are correct. We support DNSSEC Lookaside validation. DLV is used when the parent is unsigned. So, for example .com for I guess another few weeks, and some ccTLDs such as .ca. We automate the generated DLV record submission to ISCs registry.

Our monitoring module permits us to tell you if your domain will expire. So maybe you didn't notice your domain was up for renewal this month, some big companies have been known to forget. So we alert you before your signature expires, not when your domain is failed; it's too late then. We spread out the expiry of your signatures; that way as the expiry of some of your domain signatures expire – the expiry of expire? That's what happens when someone else writes this text – we don't have to resign all the records at the same time.

And the window rolls forward and the pieces get signed as they each slowly expire. This is great for large zones. It's also essential for dynamic zones such as DHCP based zones. For sites that have large amounts of zones and records, expiry spreading is rather vital as it prevents periodical massive peak loads as all the keys would expire and need to be resigned at the same time. In effect you're dossing yourself. Also with signatures that all expire at the same time, if you're running late on resigning potentially your DNS traffic load would be massive as the entire world we refresh their cache of your information at the same time again.

Now when things can go bad you need to know right away. So, for example you could have something preventing updates to your public DNS server and DNSX Signer will detect that you should have been resigned by now. So, our early warning system detects this and notifies you. Slide five please.

DNSX Resolver is part of this family of products as well. DNSX Resolver is a full DNSSEC validating DNS server. DNSSEC Lookaside validation is integrated and, as I said, implements signed domains inside unsigned TLDs. To make sure that non availability of your uplink doesn't break your own corporations DNS, visa vie your internal network, our Resolver permits you to load your own corporate trust anchors. This prevents the trust chain validation of sending out requests for .org for example, while your uplink is down. Without trust anchors your internal network would be sent serve failed responses to valid internal domain lookups.

We have a full testing mode to snap the whole state of your DNSSEC cache in the resolver, as well as the whole validating chain, to a diagnostics file for later analysis. This is great for debugging deployment. A few tricks are including in the product to help harden regular DNS traffic. To prevent response spoofing or cache poisoning queries are done towards two different name servers and then compared.

Also DNS source ports are randomized to reduce the risk of an attacker correctly responding to a request coming from a static port, such as 53. The resolver keeps rolling statistics on access per second query types – this permits you, at a glance, to judge the health of your DNS information flow. My little joke on that is it also makes it management compatible because it has nice graphics.

Lastly, finding chain of trust problems is made easy with our diagnostic tools. It chases the signature for each level of the chain, analyzes them for problems and statuses and sends you reports. Alright – I think I'm okay for time. The next three slides are our product in action. So slide six – we've selected three slides – three views of our product. I'll bring them up on here so I can tell you a little bit of what's going on. The first one is used to add a domain.

So, on the left side we see a list of all of the domains we already have added into the system. You'll see that you're not set with one particular IXFR – sorry, source DNS as well as destination DNS. So your hidden master and your authoritative server is set per zone.

So if you have multiple authoritative servers internally you can use all of them without any problems. On the right side you'll see all the parameters that we accept when you're adding a zone or a group of zones.

Here we're adding three zones at one time. So if you have a couple of thousand zones you want to add, well cut and paste, paste them there, press on "Add IXFR", as long as they grab from the same server and send to the same server then you're all set; you're all done, you can go and have lunch. As you can see we support TSIG both for receiving from the authoritative server as well as for communicating with the destination hosting server for DNS. Different key algorithms – I didn't explode them here, you can check all that in our product brochures. Of course we support NSEC3 salts as well.

Next slide please. So, the second slide shows you all of your domain statuses. The most important part here is the health as well as the message – so the state. So you have a good idea of what's going on with all of your domains at one glance – and this is live. Well not this particular slide of course, but I mean the product itself shows live.

So, in this case we see that some domains are secure but they are in need of a rollover. Some of the domains are missing DS records at their parent registrars. Most are signed, some are secure. Now the difference between signed and secure is that we have, in the case of secure we have proper parent DS relationships. I'm running out of time so I'll quickly go to the third slide.

This is the meat of it. Here is one domain selected and exploded for you so we can see everything that's going on with the domain. A trusted key at the bottom – we see all the keys that are active, their states – we only see three right now; if we were in a rollover we'd see at least one more. We see the serial numbers that are present at our origin name server as well as the destination name server where we will serve our signed zone.

If we detect in this page that the unsigned part has been update at our upstream we will show an error here and show that the signature has not yet been pushed over to – or is not valid in our destination name server. Here you can do all the manual checks that you wish for your SOA record zones, consistency is an important one. It goes up and down the trust chain and gives you a full report of what's going on at each level.

Our new version of the product again will push up the DS record directly into your registrar for you and those two lines are missing – you're seeing the old version here of the product, but those would be seen on this screen as well. That's given you a really, really warp six factor view of our product. I hope it's enlightened you somewhat about what we have to offer. And I want to thank you again for the invite.

Russ Mundy: Well thank you Patrick. And now I will just stick with Andrew – I'm sorry, I don't even want to try your last name.

Andrew Steingruebl: That's alright. So, I've been coming to ICANN for a couple of years now and I always get the question "why are you here; you're different from the other folks that are here; you're a large ecommerce site; you're a large site not an infrastructure provider, not whatever' – so, I think the flavor of just the short presentation gives you – and I've got a few notes on what we think about DNSSEC deployment and why we think it's important. You can go to the next slide.

So, if you saw Josh Powers talk at the OARC meeting he gave a lot more details about what's going on from us in the DNS world. I don't want to rehash what he said. It's a fairly simple plan. We're fairly risk averse on our DNSSEC deployment in that we really like things not to break. So, we've got several hundred zones we'll probably sign some test ones and then we'll sing several of our lower traffic ones to make sure that everything is working. The problem of course is that some of those lower traffic zones also don't get the same traffic distribution geographically or from some of our large customers, merchant partners, etc. So we're going to be pretty careful about rolling out signatures on PayPal.com. And I guess you can hop to the next slide.

When we switched over from hosting DNS entirely ourselves towards using a public – using a hosting provider, just the presence of quaday records on the responses for those name servers themselves actually caused a – some number, I won't say a whole bunch, but some number of our large customers to actually fail to do DNS resolution for us for some period of time because they past the 512-byte boundary.

And I'm glad Patrick pointed out that the PICs code's been fixed for a while. It appears that a lot of people still haven't updated so chasing those problems down is always a whole lot of fun. And as you saw from the previous SSAC report, that's probably what two or three years ago now, that the client analysis was done on home gateways and how well they're passing pacitons on – we're a little bit scared I suppose about exactly what the impacts going to be when we go push this live.

And there's a few other concerns that we're working on internally. Things like how do we want to handle key management? Are the keys for DNSSEC the same as the keys you would use to encrypt your most valuable secrets or are they more like the keys you put on a web server for doing SSL? Depending on how long you want the validity period they're maybe not that that different from SSL certificates. But you can also do a whole lot more from a denial of service standpoint if you compromise DNS keys, potentially, then if you compromise certificates. So we're still working through some of our thinking on that.

My guess is HSM is overkill for individual zone signing. But the couple of key problems surface as well, and/or things that we want to achieve. I don't want to steal some of I think what Dave and Warren are going to talk about in a little bit about uses for DNSSEC, but our motivation for deploying is we were one of the earliest adopters of EV certificates. We went to every security technology that's customer-centric, and internal as well, but customer-centric are the kinds of things that we go and deploy.

So we were early adopters. We've almost always been SSL in our websites. There were a few pages until recently that weren't over https; we've modified pretty much everything on the sore site to be https only. We were one of the supporters of the strict transport security spec, which Brian mentioned, so that we can try and prevent SSL hijacking type of attacks by hinting to clients that they should only connect with us over SSL.

And there are other security policies that we want to be able to deliver via the DNS and we want to be able to stop attacks again to the DNS. Things like what ISPs were doing with their add networks of spoofing NX domain responses on something like ww.paypal.com. I'd rather not have to put a wild card in my DNS to solve that problem.

I think that's a really inelegant solution. DNSSEC is a way for us to stop that kind of spoofing behavior, especially as it was quite harmful when one of the ad providers had a cross site scripting attack possible on the re-directory itself. So the webpage you got to on ww.paypal.com could actually be access house attacked; not a whole lot of fun to have injection on your domains.

So it's not a silver bullet; we're doing a whole bunch of things in this space. DNSSEC is important. We want to push it for our customers and I guess we want t support the whole ecosystem getting that way which was why we were supportive of all the other efforts of the root and com getting signed. And I hope I saved a few minutes off the schedule.

Russ Mundy: Well that was great. Thank you Andy. And now our last speaker panel wise is Paul Wouters, speaking for the Fedora product.

Paul Wouters: Hello. So, I'm Paul Wouters and today I'm wearing my Fedora hat. I have a few different hats and sometimes no hat, but right now you're seeing a Fedora hat. I should have worn one yes, that's true. So we'll go to the next slide. For those who don't know what Fedora is – Fedora is one of the major Linux distributions that no one uses compared to the total count of people using computers. However, it forms the basis of very many Linux distributions that people do use. So apart from the obvious very strict open source, patent free operating system, free to use, there's just a few things that make Fedora unique.

One of them is that this is the basically the leading edge for the Red Hat Enterprise Servers. So they do a lot of things. they adopt early; they were the first ones to do (inaudible); most of the Linux kernel hackers actually work for Red Hat and push their first updates into Fedora. And as we will see in a later slide, we were also early adapters for DNSSEC. Basically if you're running Fedora you're running as a guinea pig for the masses. And that's good and bad as we will see. Next slide.

I've been responsible, as a Fedora packager, for many of the packages inside distribution. A lot of the nlnet labs packages, unbound, NSD, LDNES are there. Some of them, with our work on HSM code, we have looked at supporting it properly with the packages and so that required a few more. As you know, I also work with my other hat on for a vendor that's doing DNSSEC appliances. So a lot of these packaging things we do for both Fedora and our product are like helping each other. So we do a lot of development and we push that back into Fedora and vice versa for the record, so there's a lot of feedback back into for our company to use.

There are a few things that cannot go into Fedora for various reasons and those mostly involve lawyers. So in that sense for instance, the hardware security for Sun, or I guess we should say Oracle at this point; they are binary drivers that they ship for Red Hat Enterprise Linux, not so much for Fedora, but if you know a little bit of kernel compiling and fixing a few small patches than you can make it work.

So this is available if you want to build a signer and you need an HSM, but it's somewhat difficult to use; it's not as much out of the box because not many people are using it. But these are available to add on to your Fedora thing. It will never be included in Fedora itself. Another thing that I actually have not listed here is that there is a ban on anything using ECC because of various patent claims despite EGB saying it's all fine, lawyers have said that this is still a very shaky minefield.

So you will not find any ECC code in Fedora, which also means that for instance the DNNSEC ghost algorithm is not supported in anything that's in Fedora. As a company, we're thinking of making separate packages for those people who want to test and use ghost.

Obviously since Fedora's main browser is Firefox I won't say too much because the people on my left already said it. Firefox is very important to Fedora and the DNSSEC support there is pretty important too. Then SPARTA, also sitting on my left, they have done a lot of work to make a more in depth support for DNSSEC within the applications and they've used mostly Fedora as a base for that. So you'll see a lot of packaging on their website that you can easily recompile and use on Fedora. Next slide.

A little bit of history of deploying DNSSEC in Fedora – in March 2009 unbound and bind shipped with the TLD trust anchors for DNSSEC and DLV. This was before the root was signed so what we did was we shipped a lot of these ccTLD keys and we used the Fedora update system to securely update them on your machine. We weren't sure how many people were using them because you only use this when you install a name server. But as we found out in about February 2010 when I was woken up in the middle of the night by someone from RIPE and cc, that there were in fact a lot of people using DNSSEC validation and probably without knowing it.

They just installed the name server and they just got the validation for free. Unfortunately one of the packages in one of the branches of Fedora, the keys were stuck, the package update was in the pending queue and the keys expired and this triggered a bug in bind. If you want to know more Google for "rollover or die incident" and you'll find the details on it. So my apologies for doing this – then again, we had a really good test case early on to find this bug before the root got signed so it was actually really good.

Then in December 2010 we changed all of that. The root was signed; we pulled out all the trust anchor management, we didn't need to do that anymore and at this point there's only two keys shipped – the DLV key and the root key. And per default these are enabled too, so if you install these name servers you'll get DNS validation for free. The amount of bug reporting that has happened, mostly through the Bugzilla, has actually been very few. There's been a few people complaining and usually this was a result of other intermediaries – packet fragmentation and other things – but there haven't been like one or two bugs; it's been surprisingly calm. We were actually pretty surprised by that. Next slide.

How has Fedora itself used DNSSEC and the domains have been signed since a long time – March 12 – and it means anyone using updates or if you type yum update on your Fedora install, you are secured by DNSSEC. All the domains are secured and you know for sure you are going to the right domain. DS record will be there soon. How were signing it ourselves – they're signed using a custom DNSSEC sign zone script and it's all open source and available if you want except for of course the private keys. And there is mostly two name servers and they do bind views to use GEO based IPs to help people distribute the load a little bit. And that's all working pretty well.

This is probably the most interesting slide – sorry, next slide. What do we have planned for the future? Well, we plan to put validation on every single Fedora machine. However, we do not plan to bypass all the caches that are out there. So we really want to integrate, using NetworkManager, to make sure that we're using the forwarder we get from the ISP and trying to use the cache to get all these validated answers, but we do want to have validation and security at the end point. A problem right now is that most people feel most comfortable using bind as default name server, but bind doesn't really allow for an easy update of the forwarder obtained IP addresses for their resolver. So, we're still somewhat in a designing phase in thinking how we're going to proceed. But don't worry, we're not going to do another [Denato] service attack hopefully.

Other things we're looking at and people who know my name see that I post way too many postings in the DANE newsgroup on using DNSSEC to validate SSL certificates without using certificate authorities. We're pretty much there – IETF has to go through a few more draft revisions and updates, but most of it has been, it's pretty ready to start deploying and we'll have a tool soon that will generate all these records for you and then you can start using DNSSEC for validations.

OpenDNSSEC is one package that is still missing in the repositories. I am actually working on that – there's just a few dependencies that need to be fixed before that can be pushed in, but that will be available in Fedora as well. And then one last line in my slide – this is more meant to the Fedora SSH package then to anyone in the audience. I've been asking for about three years now to enable the default SSH configuration to make use of this SSH FB records that are secured by DNSSEC. Unfortunately they are still waiting until upstream changes to default to it. So every time that I install a new Fedora machine I have to enable this validation of keys for SSH FB records and would really like this default to change. Thank you.

Russ Mundy: Okay, thank you Paul. Thank you to all the panelists. We have a few minutes built in for questions and we have mics in the room. So please come to the mic if you do have a question because it is being recorded and sent out for remote participation. But while folks are coming to the mic I do have one quick question that I will let anyone who wants to raise their hand that wants to answer – how visible do you think DNSSC should be to your customers; to your users? Should it be hidden below the sheets where nobody every sees it or should it be upfront and very visible or someplace in between? Comments/thoughts?

[background conversation]

Julie Hedlund: Please do use the microphone if you have comments or it won't be picked up for the folks listening online.

Andrew Steingruebl: Sure. To the point that I think Brian made earlier about security warnings – the problem with them as we've seen over time is that users either ignore them, don't know what they mean, or we ask them inappropriate questions. So I think should users know about DNSSEC – probably not. Should the product tell them when there's been some sort of security failure and failsafe for some definition of that word – absolutely. What does it mean to show them a DNSSEC failure if for example TLSA happens and the user looks up a key and then they connect to the website and get the wrong answer – they get a certificate mismatch – how do we evidence that to them? Do we evidence it as a DNSSEC failure or is that some other kind of failure?

So I think it's a slightly tricky question as to what that means, but in general the answer is no. Because every time you give that warning and/or a choice to the user, you're giving them a chance to make the wrong mistake and most users are going to – even if it's 50/50 – then half of the users are going to make the wrong mistake.

And on that point, I disagree with whoever said it, and maybe it was Patrick, that client access to DNSSEC records is actually really critical because not everybody is sitting on a corporate network; sometimes you're sitting on this network or sometimes you're in a coffee shop and do you trust that resolver? So it's not do you trust your corporate resolver it's do you trust whoever (inaudible) happened to hand you - and so true end on that really does actually matter in a lot of threat models.

Russ Mundy: I think Patrick had is hand up next here.

Patrick Naubert: Yes sorry. My thoughts on this are that for a major part of the population the concept of an IP address is totally alien. By extension the concept of DNS is of no importance to them except for the fact that they're typing recognizable words. People have a relationship with their computers. As techies we have a relationship with the innards and the inner working of these systems but for most of the population – most of the mortals out there have a relationship with their computers and that's it. so if we try to be transparent in pushing error messages and whatnot because we want to be full transparency up of the stack, this will only enable confusion a lot of times between the person and his relationship with the computer.

So, it's really up to the end application, to the UI, to make sense of the errors that are being pushed back up. And it's up to the UI to present the alert, to present the fact of insecurity in a fashion that can actually make sense to the user. It's up to the user to then take a decision from that. But being told that DNSSEC has determined that this particular chain of trust is incorrect and therefore you shouldn't trust this site – well we've seen viruses, pop ups that have been more congenial in their error messages to us to make sure that we click on them to get rid of such an error. So we have to be very careful, as techies, about what we make available to the end user experience.

David Lawrence: Its' David Lawrence from Akamai Technologies. At the beginning of my presentation I explained who we were but probably 95% of the people in this room already know who we were, even though 95% of the internet uses us every day most of those people don't know who we are. And we're actually okay with that. We want things to be as transparent as possible.

So we want the internet user to have a lot of confidence in the security of the overall system, but I like I think most good security is always this tradeoff between convenience and security and for the most part you want the security to be as unintrusive as possible. And so to that extent I think yes, let people know that things are secure so they can have the confidence in them, but by and large I think having it all happen behind the scenes is actually a pretty reasonable way to go.

Russ Mundy: So let’s go to the mic for questions – state your name and affiliation please.

Wes Hardaker: This is Wes Hardaker from SPARTA and my questions actually relate to user interface too. I mean a number of you have made statements to the effect that we need to roll out slowly and do things slowly in order to make sure that nothing breaks, but I’d argue that really the right way to do it is to roll it out in pieces instead of the off/on kind of attitude that I tended to hear a little bit.

So let me give you an example of some positives – if you roll out DNSSEC checking in Firefox and you pop up this big window that says you hit some DNSSEC error, here’s this really cryptic message, I think everybody here agrees that’s not going to work; the end user is going to be clueless. At the same time, having the support in Firefox and having the support in routers and having support in SSH to check fingerprints, but maybe having it off by default for a while would be great so people actually can turn it on.

For example, I actually was one of the people that wrote the patch to Firefox in order to get DNS in it and we had to mess with a lot of low level code, but in the end there's still a check box in the options menu that says do you want to enable DNSSEC or not. And if you turn it on then you can, but users don't need to necessarily need to know how to turn it on and off unless they go into the advanced section.

But the reason that I'm suggesting turning it on at last in slow pieces is that there is a lot of advantages to it that can help the user interface. And my favorite example is actually also Firefox related where if you type in an unknown address into Firefox today you get this four bullet list of all the ways that you might have messed up. And it asks you very generic questions – it says sis you mistype the name; is your network on; is your computer plugged in – I mean they have no idea because you're clueless. All Firefox got back was this I didn't get an answer so the only thing they can present to the user is I don't know what went wrong.

Well, with DNSSEC actually on, one of the valuable things we came up with is we were able to shorten that because if you can prove through DNSSEC that you got an answer and that record does not exist, you can shorten that bullet list to "the place you're trying to go doesn't exist". You still may have made a typo, but it wasn’t the network, it wasn’t your system administrator, it wasn’t anything else – I’ve proved positively you’re trying to go to a place that doesn’t exist.

So there’s a lot of advantages to actually pushing out DNSSEC and getting some of these low lying things – getting SSH fingerprints verified underneath; getting SSL turned on by default without ever going through http – all that kind of stuff can be hidden underneath the hood and can be turned on without necessarily having to go “how do we display the error to the end user.”

Russ Mundy: Oh, sure Patrick.

Patrick Naubert: Sorry. It’s a great opportunity for me to actually stand on my soapbox. A big fan of Firefox… One of the things that makes total sense in phasing in security is to be able to give the choice to the user. The biggest problem that I see is again, with the relationship with your computer. We’re used to hunting – again, us techies are used to hunting for these parameters that would increase security and not downgrade our experience with our machines. The problem again, agreeing with you, is one of interface. Your view of it is product-centric.

I would like to argue that this has to shift somewhat to operating system centric where the user would be presented with all the possible things that they could do to increase security and then they get to choose it. But if that is not there, then the trap becomes that the user will have a false sense of security as he enables DNSSEC, for example, in Firefox and believes that it's active in other portions of the operating system, maybe they don't know better.

It depends again on how at the UI level from Firefox you're presenting the opportunity to increase security. Are you going to say well make more secure and you have a slide bar that goes from none to top and you take care of DNSSEC in the background while suddenly the user is under the impression that this applies to the whole operating system.

Or as if you have a checkbox that says "enable DNSSEC" – the user is clueless, doesn't really know what that means – heck I'll do it or I won't do it. So it's good to give a choice to the user but at the same time if the UI does not present some sort of enlightenment to the user for that choice, the choice is bogus and really has no value.

Dave Crocker: Dave Crocker; Brandenburg InternetWorking. I'm just delighted that Russ opened with this question and the conversation is going where it's going. The beginning of my career was focused on human computer interaction – words these days used for that are usability and user experience design. I got out of that because protocols are much easier to design then good user interfaces. And it's also much easier to get funding to do that work.

And unfortunately the kind of comments in the presentations and being made so far represent what I think is an industry standard display of dissociative identity disorder, which is a small pun in this particular context, that used to be called multiple personality. Many of the comments are showing an awful lot of insight about usability issues, but deep down there's this other personality that is consistent in the community and in the comments that we need to inform users of stuff; that we need to make uses responsible, hold them responsible for stuff.

All of the experience to date for anything in the mass market says that's fundamentally flawed. You cannot expect – and we've had some comments that reflect the kind of thing I'm saying right now, but overall including from the same people, the comments go at odds. And that is we need the infrastructure services to be able to take advantage of this stuff and take advantage of it well. The instant you talk about presenting anything interesting to the user – choice or problems or anything like that – let me suggest that you just step out of what is an area of expertise and into an area of ignorance. And the really sad part is it's not just your ignorance, it's the ignorance of the experts.

There is a conference every year called SOUPS, which is about usable security, and let me suggest to you it's enlightening to go to it because it's really depressing. The state of the art is very poor. And what I'm suggesting in the anti-abuse arena is every time somebody says we need to give the users more choice or more information, they need to be condemned.

Luke Stans: So, my name is Luke Stans; I work at Mozilla running random security things. I've actually been to SOUP – it is sometimes really interesting and sometimes a little heavy on the analysis and a little light on the conclusions but. So I'm a believer in minimizing interaction around security with the user. I think generally speaking it's sort of kicking the can down the road and having the person least qualified being forced to make a decision that will affect them directly if they make the wrong decision. That said, you can't avoid it entirely.

At some point if there is something wrong with their site, the user has to make a choice. You can take the choice away entirely, which is sort of what we did in Firefox. We sort of for the most part said this site is broken, please try again later. And if you're determined you can make an exception. The reality is if you provide no exception mechanism the user will just go use a different browser and make the same bad decision they would make anyway. That is a choice you're presenting to a user.

So, I think that we should focus on putting in the plumbing to support DNSSEC. I think there are some really great examples where it could improve DNS warnings, and we could provide HSTS over DNSSEC and probably a lot of other valuable services. I don't think I'd ever want to present a DNSSEC specific UI. I think I want to figure out how to – the goal is to inform the user whether or not the connection or the interaction that they're having with the website is trustworthy or not. I agree DNSSEC can help do that – great. It shouldn't internally conflict with the signals because that's going to be counterproductive.

Russ Mundy: We're actually a little bit over time so if we could have a 30 second or so limit on the questions and answers we'll try to move through. And I think that mic was next.

Dave Piscitello: Yeah, this is Dave Piscitello from ICANN. It's interesting Dave Crocker and I almost had this exact same experience three weeks ago at MAWG; we were talking about relatively the same problem. I think if the answer to Dave saying if you want to try to present things to the user and then you should be condemned is that if not presenting something informative or assisting in resolving to the user then what – what's the alternative? The user will almost invariably do something – he's not going to stop. So, I'd suggest that maybe it would be worthwhile for the DNSSEC community to consider what the Anti-Phishing Working Group did with their phishing landing redirection page.

Instead of just trying to have engineers who are button people and who are much more analytic then the general public do this, what they attempted to do was to go out through Lori Craner at CNU and through user groups and focus groups with various pages that actually represented some sort of education about the fact that they just had nearly been phished.

And so you can go and you can look at the page and you can look at some of the papers that Lori's written and maybe what you want to do is collectively come up with some notion that is similar to a redirection page that helps people go to a trusted page, one that actually does DNSSEC resolve with an appropriate signature, and gives you some information that tells you what went wrong and how you might be able to rectify it. Thank you.

Russ Mundy: Jim, go ahead.

Jim Galvin: Thank you. Jim Galvin from Afilias. I want to first say that I agree with everything that Dave Crocker said, that was an excellent comment. And I want to make a simple statement that agrees with that with a slightly different foundation. I've said this before and I'll continue this from time to time – with security there is very little room for gradation. I mean if you step back and you think about security from that sense it's either on or it's off; you're either secure or you're not; or whatever definition of security that you're working with.

And in that context if SSL has taught us nothing the one lesson we should have learned is you cannot give the user the choice. If you give them a choice they're going to flip through. It's as simple as that because they're not going to think about it, they don't know enough to think about it, they have no way to judge the options. We need an entirely different means of educating users to judge options if you're going to do that to them.

So my simple statement about DNSSEC is it's either signed or it's not – and if it's signed it's okay and if it's not it simply doesn't exist. And that's an ideal world that we need to get to. Now I realize there's some consequences about that in a transition period, but that's the space we want to be in. If it's not signed it's simply not there.

Andrew Steingruebl: I'll make one quick comment on that. We've recently seen a huge push by a bunch of social networking sites to enable https everywhere – like Twitter just announced that they were doing it and supporting it universally and added user interface to it yesterday – and I think that, I think it was Brian who said it earlier that from the Firefox point of view if you're trying to set up a, if you're doing a DNS lookup and you're not trying to also set up a secure channel, then the value proposition is a bit suspect. Right?

So none of these are a silver bullet and while I'm all for the universal deployment of that let's not stop there and say that we've actually really accomplished something if we get there. Because if you're still speaking now, if you're talking to the right IP address but it's being intercepted in the middle by somebody, that communication is being intercepted in the middle, you haven't really accomplished anything. So it's not one or the other, it's both and a whole bunch of other things added into the mix.

Russ Mundy: Okay, one last…

Bill Smith: Sure, I'm Bill Smith with PayPal. In full disclosure I work for Andy. Exactly. The point I wanted to make following on some of the more recent comments is we need; basically we need to move from a world where security is an option in a sense and into a point where it's a requirement. And either you don't get to turn it off or turning it off is really hard. And the other thing we have to do is move from, I think, we need to change the way we as a community think. And basically again, turn security on by default, but do so in meaningful ways. This is not intended as a comment against any company here, but we shouldn't have wireless routers that have the user name as admin and the password is admin by default.

That's just – I know, and it's mine, too, and…Quick story – going back when I worked for a lottery company, we introduced these slips where you could mark your bet that you wanted to make on the lotto game and we gave an example on the back of the slip. What do you think the most frequently occurring bet was? Because people saw it and they said oh, that's how I'm supposed to mark the slip and did exactly what they were told.

Anyway, we need to turn stuff on by default but we shouldn't allow them to say oh I'll just keep the password the way it is.

Don’t set a password or make it something crazy – I don’t know the solution, but we need to change the way we think, I believe. And this is going to take time, but if we don’t start we’ll never get anywhere. Thanks.

Russ Mundy: Okay. Thank you very much everybody. We do have – one thing that as folks that have been involved with the .gov activity, that I wanted to ask David to comment on because one of the reasons we’re hearing .gov uptake isn’t as much as kind of the inverse of where you were saying there is no demand – what we’ve been hearing from the .gov participants and people that are trying to deploy this, that they aren’t going to deploy until Akamai provides a broader set of support. Does that…?

David Lawrence: I’m happy to address that one because we, before coming here we wanted to understand from the government sector services people what they were hearing and that’s still what they’re not getting that message. So, somewhere there is a drop in the communication chain. The people that are feeling like that apparently have to tell the people that are talking to our people that that’s the message they want to communicate. And personally, as an engineer who’s been involved with DNSSEC for more than a decade now, I’d be happy to have the additional government push. But we’re a business and we prioritize things based on what the customers are demanding and we’re not hearing from the government.

Russ Mundy: Okay, well that’s good. I appreciate that. And this is another reason that I wanted to ask that is because for those that want to have DNSSEC and want to have it available, you’ve got to ask the right place, the right organizations at the right time. And too many times we end up, a bunch of techies talking to each other and don’t get the marketing people involved.

And that’s sometimes what it ends up taking is your buyers, your contractors, your marketing people – if you want to press forward with it, you need to get those folks involved also. And I think we’ve – Steve did you? We have, we’ve eaten a lot of extra time here with a little bit of the A/V stuff, but I think we’re going to jump right to our next activity and shorten our break – break is gone…That’s very short. Okay. Okay, thank you panelist.

Steve Crocker: Alright. We’re going to keep the pace up here a bit. I’m pleased to introduce my brother Dave Crocker. Historically I’ve focused quite a lot on security and he’s focused on mail and on user level things, but over the years, which are now adding up rather extensively, every once in a while we find ourselves intersecting and even on occasion in role reversal situations. So this is one of those pleasurable opportunities where our paths have intersected and so I’m very pleased to have Dave talking about the intersection and the fusion between DNSSEC and higher level applications.

Dave Crocker: Thank you Steve. HI folks. I apologize for the coughing.

Steve Crocker: Oh and keep it short.

Dave Crocker: I've been hearing that all my life. So rather than saying your break is gone let's think of this as your break. As Steve says, I don't do security except every now and then I find myself in it. for the last six years or seven years I've been working on DKIM. And if you have any questions about algorithms I'm the wrong guy; I don't know or care about them. I care about the labels and I care about the security telling me, security community telling me that a particular algorithm works or how it works and what it does but not the underlying part. What I'm interested in is the assemblage of algorithms into something that is a service. What I've noticed in trying to interact with the security community is some underlying problems.

Everybody I've dealt with in the security field knows a lot and is careful about what they say, but the consistencies from person to person are highly variable, which is to say things aren't very consistent. And so for expel, I'm starting to believe that we need to ban the word security from all but a few specialized discussions because every time the word gets used in the world that I walk around in it's used as an umbrella term rather than a specific term and people walk away thinking that they understand what they're being protected from but they don't. So TLS gives us security – okay, my data is now safe, I know who I'm talking to, they know who I am – oops, only one of that might be right and only a little bit.

And similarly, it turns out I've been discovering even more specialized words like authentication. So with DKIM, the first version we did, the signing specification, basically tells you that we protect the message. Well in fact, that's not what DKIM does. DKIM protects a name that is associated with a message and it isn't the "from" field.

The reason I want to harp on that is what we I believe need in the more general world of usable security at the application level is a variety of different purpose built security services in which we're very careful about what we say it does and what we provide to the consumers of it – they may or may not be users. And so down at the bottom take a look at the three bullets.

The first one was a definition that I was quite pleased with in that it was very careful for example, to distinguish between integrity and authentication and two different kinds of authentication. That's more carefulness than I'm used to seeing in summary descriptions of a security service capability. The second bullet was the string that we have in the original DKIM signing spec for what DKIM does – the second bullet is what we currently say.

And I will tell you I think these are fundamentally different and that the second bullet – the first DKIM bullet – is wrong, we've been misleading people because we didn't adequately think through what it was we're doing. So DKIM comes out of some work at Yahoo which was very clever. It's incredibly ugly in some ways because it starts with an installed base of services like the DNS and lays on top of it some capabilities.

And it has to be rather creative in the way it does that, but the core benefit of that capability is it gives you self-certifying keys associated with an identifier – it uses the DNS to do this; it does it in a way that's practical and now highly proved. I don't mean proof in a math sense, I mean it's pragmatic.

It packages the per unit – originally per message – information, the signature, is a way that is out of band of user visibility. And it does that by putting it in email into a header field. So it's there with the message but it's not intruding on the end user. That's really important if you're worrying about adoption and support. Burdening the user the way that SMIME or PGP does is very intrusive and problematic.

And the other thing it does is it has some algorithms for allowing the signature to survive certain kinds of perturbations which are deemed benign. And those canonicalization algorithms are useful, interesting, and minimal at the current point but one can plug in other algorithms for that carefully. So there is a balance between the robustness against manipulations that applications sometimes experience in transit versus sufficient security enforcement.

What Murray Kucherawy, who is sitting over there, and I did recently was to extract out of DKIM the core capabilities that DKIM does that are not specific to the DKIM service. They are generic security related services. In other words, it creates a library. And with that library it also provides a generic template for doing authentication so that if somebody wants to create a purpose built application security service they only have the hard part to do.

They don't have any of the grunge work. They don't have to create a key service. They don't have to define any of those mechanisms. If they can live with the existing algorithms hashing and cryptoalgorithms, they can use those. And they have to worry about, for example, what does this service really mean; what specifically does the presence of a signature mean. Does it mean that the "from" field is valid? Does it mean that the contents are valid? Does it mean that the attached identifier touched the message but isn't making any statements about the validity of the message?

That three – range of three I gave you are examples of the different things DKIM might do of which it only does one, but similarly, we can apply similar kinds of choices for other applications. This reduces down to this table that seems to be – and I say seems to be because this work is only a few weeks old – seems to be a way of parameterizing basic security services at an application level.

So, the first one is how do I know – sorry – how do I take the security information and link it to the data; what's the mechanism for doing that, that's a mechanical point. The second is how do I inform the receiving side; whether it's a user, I hope not, or a receiving application, I hope or an infrastructure service, I hope even more – how do I indicate that this mechanism is in force. Third one is if it's in force what does it mean. And the third one is for the authentication template that's in the spec – how do I map that template to a specific, such as email or a MIME object or a VCard or VCAL object or an XML object. This is a chronic cough, I apologize.

The first instance that we did was, and it's in the base spec for DOSETA, is to redo DKIM in terms of DOSETA. That's in the appendix of the base spec. The second thing we did was to do MIME authentication in this method; and you'll see that this defines a content authentication field, I'm pretty sure one has not previously been defined from the searching I did. So the presence of that MIME header field says MIME auth is in effect. The semantics – it turns out I don't have agreement yet on what the semantics should mean and I started to come with a meaning then talked to a few people and I decided I really didn't know what the semantics should be, what's the most useful.

As we try to get people interested in doing MIME authentication using this mechanism, we should get some agreement on what the authentication signature means. And thirdly was the mapping, the template that you'll find in the spec, defines a header content kind of model, standard two part model that occurs in many situations. So it maps between the MIME header and the MIME body or body part and that's fairly straightforward.

So, what does any of this have to do with DNSSEC and at one level the answer is nothing, but that wouldn't be a good position to take in this group and it wasn't a motivation for adding me to the presentation. The answer is I think that security for infrastructure services and security for applications really is a different beast. For one thing we need lots of different security services for applications.

For another, the entire world of operations and administration is completely different for the infrastructure services than it is for applications. The folks who need to be able to manipulate things for applications usually are different staff with different skill sets than are doing the infrastructure services. And the folks who are doing infrastructure services typically are kind of protective of making changes to things – as well they should be.

So separating these functions and making it easier to have lots of different application services without touching the underlying service seems to me a very good idea, but there needs to be some complimentary operation here so that there is an end to end set of security assertions that can be made. And for that, I think having DNSSEC be able to guarantee that whatever the data are in the

DNS, they were put there by the owner – might not actually be valid because the owner might screw up, but you can be sure that no one else put them there.

That's an incredibly powerful statement as you get more and more application services that are depending on the validity of the end to end exchange. Right now we know there's a whole – that the underlying mechanism can be attacked and taken over; DNSSEC fixes that. There is not a lot of end organization market pull for this because the applications that they're using don't see the need. What we need is to fish around – that's a small pun – for applications that excite end organizations and businesses where the risk of failure, the risk of compromise is sufficiently high, that they will be really excited to have an end to end assertion that is strong and safe.

Ultimately therefore this presentation is a solicitation for interest in the topic and an effort to proceed forward looking for those applications. The number of different things that DOSETA can be used for I've tried to indicate both in terms of the security service and in terms of the different kinds of object, the model is very generic and I think that it can be applied easily in lots of places. The easily is the interesting part here. DOSETA makes the mechanics of creating and deploying a security service fundamentally easier than we've been having to date. Thank you. Questions/comments?

Russ Mundy: Thank you. This is very unusual to let Dave Crocker off the hook without further questions, but thank you Dave. Now, Warren and Jay. One quick question, we're just about on time, but let's hear a very quick question here.

Ben Wilson: I'm a little bit confused about what DOSETA is in terms of is it like a protocol, is it a framework, where does it tie into- I mean if you start with the technical framework then what's the next step; where does it tie into the application.

Dave Crocker: Sorry about that. DOSETA is a set of protocol constructs that can be tailored to a specific security protocol. The model is DKIM, but it's made more generic so that it can be tailored. But I would class it therefore as a generic protocol that can be tailored. And it's defined as a library of security functions with a template for doing authentication, but I'm quite convinced it can also be used for other security services, notably confidentiality.

Russ Mundy: Okay, thank you. So, I think we're ready to move onto the next group. And we'll start with Warren Kumari from Google. Go ahead, Warren.

Warren Kumari: As Russ said, I'm Warren Kumari. I co-chair the DANE Working Group along with Ondrej Sury, who I think I saw somewhere in the room. And unfortunately I've only got 10 minutes for this. I think I've got a bit more content then that so I'm going to be talking very quickly. But I would like this to be a conversation so if anything I say is unclear or if you think I'm just making things up please interrupt me and let me know.

In order to talk about some of the stuff we're doing in the DANE Working Group I first need to provide some background. And for all of the SSL and TLS weenies in the room, I'm sorry I'm simplifying things hugely and glossing over all sorts of important bits. Okay, so when you connect to httpswww.example.com you're using SSL or more correctly TLS to build an encrypted connection.

In order to do that you need to get a public key and that comes shipped to you in a PKIX cert. The important bit with all of this is you need to make sure that you've got the right cert. This prevents "man in the middle" hijacking attacks – in this picture the scary devil guy at the top is intercepting a certificate from example.com and is replacing it with his own.

So, the way that you actually protect against man in the middle attacks is using PKI or public key infrastructure. What happens is example.com generates a public and private key and obviously keeps the private part separate, like you do in DNSSEC. Then it puts the public key in a certificate signing request. It ships this off to a certificate authority and the CA contacts example.com and verifies that the key that they've received is actually the one that example.com wanted to send.

Assuming this checks out correctly, they issue the certificate. And this contains the public part of the key; the host name for the certificate, where it's going to be used; and the CA signs it. And the signature actually binds the key and the host name together. So if somebody gets the cert and tries to fiddle with either part of it, the signature obviously will fail.

When you come to actually use the certificate, you , or more correctly your browser, downloads the certificate and the first thing it does is it checks to make sure that the host name in the search matches the place you're trying to connect to. It does a bunch of other validity checking like make sure this certificate hasn't expired; make sure it hasn't been issued sometime in the future; that this certificate is actually allowed to be used for securing web connections – things like that. It then checks the signature, makes sure that that's valid and assuming that everything is we go ahead and connect.

But we haven't actually solved any problem here yet. The initial issue was that we didn't have away to validate the key for example.com and now you've got the signature saying this is the key, this is example.com, but you don't have a way to validate the signature because you don't have the key for the CA yet. So, luckily, or somewhat luckily, these keys come preconfigured in your browser, in your operating system. These are sort of root keys – you can think of it sort of like the DNSSEC trust anchor.

And you inherently trust all of these configured CA certificates that come in the browser.

So OSX comes with around 163 of them – you can see a list of them there. Some of them are owned by nation states, some of them are private companies. Mozilla/Firefox – as of January 2011 comes with 155. I don't know is this number is still valid but I think it's close. Internet Explorer/Windows – it's kind of unclear how many it comes with. Initially it comes with a very small trust anchor store, but these get dynamically updated so it's hard to figure out exactly what the number is because the numbers keep changing.

So, as well as the CA certs that come preconfigured in the browser, some of the CAs have sort of intermediate signing certs. So these are certificates that have signing authority so they can also be used to attest to the binding between the key and the host name. So in total there are around 1400 of these and initially this seems really good. I mean if you're example.com and you need to get a cert more places is good, competition should drive down price. You can go with somebody you trust which all seems great; unfortunately, no.

When a relying party comes to use a certificate they've got no way of knowing ahead of time which CA should actually have signed it. so if example.com gets their certificates from Kimodo, but when they actually, when somebody downloads one they find it signed by somebody else they can trust, the relying party has no way of knowing that this isn't supposed to be happening.

And there are many reasons that a certificate could have been signed by a different CA then the one that was expected. One of the more likely is just the CA has been tricked into doing this – that's the incompetent case. Things that are a lot more scary is the CA could have become malicious and could be doing this on purpose. Another option is the compelled CA case where law enforcement goes along and says we need a certificate site for lawful intercept; or more likely we need an intermediate cert and then we can on the fly do sort of man in the middle type activities.

And there is actually hardware sold specifically for law enforcement that does this. You plug Ethernet in one side, you put a signing cert on it and on the fly it will do man in the middle hijack attacks. So, this doesn't happen all that often. If you're browsing to a website and you see the locked icon, chances are everything is okay. But when it's not okay, that's really not good. And this joke probably only works in the US.

So, what we're trying to do in the DANE Working Group is we've realized that at the root of all these problems is the fact that there are just way too many trust anchors. What's nice is DNSSEC has a single trust anchor. It's free to use, you don't really have to pay for it. It provides a way to securely publish information and then validate that information. At a specific node in the DNS only a domain owner can publish stuff. There’s an easy discovery mechanism for the information - the DNS itself provides this discovery mechanism. And it’s also got authenticated denial of existence so you can authoritatively state there is no such record here.

So what we’re actually proposing is that sites can take their existing certificates, you calculate a hash of the cert and you publish this is the DNS – I know we’ve covered the new resource record called TLSA – you publish it in the DNS and you sign it using DNSSEC. When relying parties actually want to use the cert, they get the cert in the normal way they would with TLS, they calculate the hash and then they compare it to the information published in DNS. And if everything matches, all is good. You’ve got proof that the certificate that you’ve downloaded is the one that the site intended you to have.

If on the other hand they don’t match, you’ve got very good evidence that something bad is happening. Somebody is currently doing a man in the middle hijacking attack; it’s also possible the DNS admin messed up and this record has just expired, but you should probably assume that something bas is trying to be done against you. And this already provides a lot of utility. If you’re a large company – Microsoft, Yahoo, PayPal – somebody like that, you would like to be able to give your users a way to ensure that the PayPal that they’ve got to is the correct PayPal. So we think that there’s already is a big one. But we’ve got more.

In order to get a certificate for a domain, all you really need to do these days is prove that you control the domain. Normally this works by the CA sends an authentication token to an email address at the domain name – like postmaster@ or hostmaster@ - and then the site administrator gets this toke, ships it back to the CA and that sort of proves that they're in control of the domain.

The issue is anybody who controls the DNS for a domain can point the mail servers to wherever they like. So if you control the DNS for a domain you have the ability to actually get a certificate for the domain. So a rogue DSN admin can get a certificate for any of the domains that are under his control.

So, what the CAs actually do is they bind the key and the host name together. They're sort of attesting that this binding exists; this is the correct key for this host name. What we think is if you generate your own self-signed certificate and you publish this in the DNS and you know, in a TLSA record and you've signed it, you can self attest to the fact that this is a correct cert.

And seeing as though the only person who can actually publish information in the DNS is the DNS admin, it doesn't really open up additional avenues for attackers to fiddle with certs. So we think that if you publish a self-signed cert and it's signed using this, you have about the same level of trustworthiness in the site as if it was a CA signed cert; and some people think potentially a bunch more trustworthiness.

So you know, DNSSEC was originally sold as a means to prevent spoofing and cache poisoning and things like that, but it's actually a secure publishing method and it automatically limits where a user can publish information; you can only publish information at your node in the tree or stuff underneath you.

So we think that this opens the door for all sorts of interesting new applications. We think DANE is a cool and exciting thing. Jay is going to talk about a bunch more new and interesting ones. If you want any more info on this, either find myself or Ondrej Sury, who I don't actually see in the room, otherwise our charter is over there if you'd like to be involved that would be great. And I think that's the end.

Russ Mundy: Thank you Warren. Now let's just jump right to Jay.

Jay Daley: Okay, thank you. My name is Jay Daley from .nz. First of all, a little disclaimer – a short one. As you'd expect this is highly opinionated; I may be wrong. Some other bits, I may not actually understand what I'm talking about and as this is entirely futurology none of this may happen ever. Thank you. Next slide.

This is about what may happen in DNS and applications as a result of DNSSEC. It's something I think a number of us have been thinking about for a number of years, largely to persuade people that there was more to DNSSEC rather than just fixing a problem; that it has new opportunities as well to come from it.

Within DNS then we already do address mapping, that's fairly clear, that's well understood. And what DANE is doing, through the use of DNSSEC, is significantly enhancing the way that DNS can provide security for people for applications. And there are – let's be clear about this – there are 205 million domain names; a large portion of those run websites. And approximately 2% of those websites are protected by valid signed third party certificates by a certificate authority. If 2% is anybody's definition of success then please put your hand up; it certainly is not mine. If the goal is for us to have secure web access, then the certificate authority mechanism has failed and failed dramatically.

And what DANE hopefully will do is fix that problem and we will end up with ubiquitous https security. So, if it wasn't clear from Warren's presentation, the good thing about DANE is it will be free. You can create your own certificates and publish them in the DNS. And then you can share those certificates across multiple hosts because you don't have to have the link between the domain name and the host that's running it. So you can deploy the proper infrastructure in a lot of flexible ways.

There are some other resource records in DNS that deal with security – SSH fingerprint being one of those. Those are really unused at the moment as far as most people know. Certainly those who run data collection to look for people looking for those types externally see very little of them. Obviously most of them would be internal only within an organization, but I think we can safely say they are unused. And one of the hopeful side effects of DANE will be that those security types are used much more. We may even see the GEO types, which have been long deprecated, being used as well; wouldn't that be a bit of fun.

Okay so then policy – now this relates to Dave's talk really. By policy here I mean we are more and more seeing things that say if you connect from this IP address range you need to connect like this – or I will only talk to you if you come from this IP address range – or well, a variety of other things. Now, this I think is likely to be a growth area; its' security plus in a sense. It's how to implement the security; or possibly other bits and pieces as well.

It's making statements within the DNS that help people understand should they connect and under what circumstances should they connect and how should they do that. I personally think that squeezing all of that into one line in a resource record is madness, but that's almost certainly going to be the way that it will go.

Okay, there are a lot of useful characteristics that come with DNS. Some come by design and there may be more on this list, it's scalable, it's distributed, it roots around failures, it's compact, it's replicating, and now the big one, it's secure – not fully secure, but secure in the way that we need it to be as a directory service. There are also some that come by accident, one of which is that it's generally firewall transparent. And the other thing is that it is exceptionally profitable. There has been an enormous amount of money put into DNS and things because domain names have been such big business.

They are after all paying for what is our magnificent venue for a technical conversation. So, this is something that generates investment from manufacturers, from suppliers, and others in equipment and software and other things; all of which people look at and buy and think well how can I leverage this, how can I do more with this.

So I think more and more people are thinking now that we have security added, with all of these other characteristics, it seems natural for us to use DNS for databases much more similar to ordinary databases. And as another completely probably wrong prediction, I think we may see more requirements for clever database features needed in the DNS – point us to things, indexes, possible bits and pieces like that. but the big danger to me is that one of the fantastic characteristics of DNS which is lose synchronization, may now be under threat because databases don't generally do well with lose synchronization.

Okay, now those of us who are TLD operators generally hold our fate in our own hands. If we're hurt then that's our fault. We haven't enough penetration tests; we haven't had a proper security cultures – all sorts of things. But there's one big thing that scares all of us operators I'm sure and anybody who says no is lying, and that's distributed denial of service attack. There is an enormous asymmetry out there. We run, even if you take somebody big, somebody very big, you might have the 40 gig worth of bandwidth coming into their DNS clusters or something like that, or possibly more. They can still be taken out by a botnet of a couple of million because the asymmetry is still so large.

For many years people have thought that peer to peer DNS infrastructure is a natural mitigation about this. We distribute the serving of DNS across as many computers as the attackers might have and therefore we no longer have the asymmetry, the bandwidth asymmetry that puts DNS under so much threat. Now, peer to peer though all about trust. And the big problem has come how you deal with bad actors within that circumstance. DNSSEC is one part of the fixing that – it provides trust for the data integrity.

So if we did have a peer to peer infrastructure than any publish of the data can at least ensure that their data is published from integrity. Now there is still some other very big bits to deal with here – there is still no trust for server integrity for example, for other bad actors who are deliberately sending broken things, but that's not too hard to solve. Reputational systems are out there and people do use those. And then getting acceptable performance out of it is not trivial either – and I'm sure someone else can do that rather than me.

Okay, so the final slide then about some other side effects. Cryptography and key management in enterprises is now common; https is a good example. Many run intranets with VPNs and have certificates for that or they have client certificates for their customers to access them with. And DNSSEC is another one of those crypto functions. But what DNSSEC brings compared to those others is some good examples of the management process.

It has embedded within it, as a good practice rather than as an absolute requirement, this split between zone signing keys and key signing keys. It has the concept of signing ceremonies. It has the concept of how you protect the key and distribute it. These are all coming out of the good practice that's been set by the way that IANA does things.

So I would hope that that makes it clear no to every enterprise, that they need a CA function entirely within that enterprise. That implements these good practices; that creates organizational root keys, uses those to create intermediate keys and then uses those for the signing other keys that are then used within the enterprise whether those are http certificates for whatever.

That they implement the best practice of signing ceremonies or something equivalent; of HSMS key storage and that types of stuff;

and they look holistically across their organization for all of their crypto needs. And then if somebody clever can sort out translating keys between all the different formats, that would make life easier as well.

So, that's it for me. Any questions?

Russ Mundy: Thank you. We have actually about five minutes for questions if folks want to ask any, otherwise we'll move to our – oh, here comes Andrew to the mic I believe. Name and affiliation please.

Andrew Sullivan: Andrew Sullivan and here I guess I'm wearing my [Shinkuro] hat. This discussion, and actually a couple of the other ones as well, have really tweaked for me the problem that it's going to be awfully hard for us to get users for all of this nice stuff that we're putting into the DNS if we don't have an easy way for applications to reach down and get that stuff and get at our info is not the answer. So unless we start solving that problem kind of in parallel to all of this, we're going to have a hard time. We're going to build all this nice stuff like DANE is doing and all of the things you've just suggested and the applications are going to great how do I get to it? Does anybody have any thoughts about what to do about that?

Jay Daley: Is the question there whether they just get it from DNS or whether they need to get the trust as well as getting it from DNS?

Andrew Sullivan: Well, for instance, it does seem to me that for a browser to use what they've just gotten out of DANE they're going to need to be able to know with absolute certainty that that answer was validated. And right now the answer to that is well, build your own resolver inside your application and I don't think that's a good answer.

Warren Kumari: Yeah, the use of DANE very much requires that the application has done its own checking or has got the answer from a name server that it has a secured connection to and can validate that. This is obviously something that needs to be solved. It seems like there might be a working group in the IETF that could deal with that.

Jay Daley: The other thing is for operating system designers to put interprocess certificates into place so that you get a trusted answer from your DNS resolution library.

Andrew Sullivan: So, I guess that's part of what I'm asking. If we're going to build portable operations here how are you going to do that? There's no handle in [Posix] for this for sure. So, is anybody working on that? I mean you guys are.

Russ Mundy: While he's going to the mic – there's also been an API published previously in a draft that needs updating. Go ahead, Paul.

Paul Wouters: Paul from the Fedora project. So obviously we solved this was by putting the resolver on every host and still trying to use the caching infrastructure. Like every application should not get a validator.

Wes Hardaker: Wes Hardaker; SPARTA. There are at least two and I believe three libraries that already provide this, the problem is that there's not a consistent interface. So Andrew, it's certainly possible because Firefox already has the code in it, from our stuff to do exactly what we're talking about. The problem is there's not standardization for something similar to get host by name or get add or info or stuff like that is consistent across all those libraries and that's what we need to fix.

Brian Smith: I saw references in I think both your presentations that DNSSEC is free – I'm kind of skeptical that it's free, especially in a high quality implementation. I think later we're going to be hearing from people who are selling services because it's so hard to do. Do you really think that it's going to be a lower cost compared to the CA system that we have now?

Warren Kumari: Well, presumably people are going to have to implement DNSSEC architecture to you know, serve their records and stuff like that. The incremental cost to do stuff like that hopefully should be fairly small. And you're doing it yourself; you're not paying it to an external company. But yes, DNSSEC is not free; the use of trust anchor is free.

Russ Mundy: Okay, thank you Warren; thank you Jay and thank you for the questions and comments from the audience. Now we'll move to our next panel presentations. Thank you.

Steve Crocker: So, one of our standard features is to focus on updates on implementations from around the region – this region, North America, is relatively few separate countries, but we have nonetheless, two very big pieces of news here. So, let me start with Jacques Latour from the Canadian Registry CIRA. Thanks.

Jacques Latour: Hello. Alright so today I'm going to talk about our plans for implementing DNSSEC. So I just want to start – this is our current state, I want to show you where we are, where we're going. Right now we just finished the migration of our registry to new EPP platform. That took about – it was an 18 month project and I guess most of our focus over the last year and a half, two years was around doing that. So, what it means is that around December of this year we started the real planning of doing DNSSEC internally. So we're still in our planning phase.

We did migrate on October 12 – our registry with a brand new hardware platform, operating system, the whole thing. We did a 24-hour cutover from our old platform to the new one. And it went well. It was a full cutover – all the registrars were using proprietary interface to EPP, so that worked out well.

So now we're working on DNSSEC. So we're building a plan as we speak trying to outline all the issues. Our key focus is not around developing the DNSSEC technology; our focus is around using what's available out there in an operation mode. The operation team that we have today is a bit immature in terms of IT processes. What DNSSEC does is it requires a lot of operation processes to implement the DNSSEC. So that's our biggest challenge. Right now it's all about research, training, getting the people to really understand what it is, what it impacts in the operation side of it, how do we support, how do we maintain – that's the biggest challenge that we're having right now.

So the plan is about doing solution architecture and design – so we need to do that on our own because we need to understand what we have, how we support, and how we manage. We need to build a detailed project plan to get there. We need to execute the project in the phase and in the proper method. Tones of process development – that's a key thing. We need to develop processes for managing the keys, for managing the operation, for managing incident management and all of that stuff. And internally we have a project management office – they have a risk management process and that register is getting filled up with a lot of issues in there.

Doing DNSSEC is not just putting in a signer, an HSM module and getting it done – it impacts your whole ecosystem for a ccTLD. So we need to get sure that the registrars support

DNSSEC, we need the EPP extension to support that, we need a database structure to support the DS record. And then there's the stuff we talk here about – the singer and HSM and all of that stuff and key management. So it's broad – it's more than just a bump in the wire; it touches all of the business, including finance.

So right now we're looking at solution architecture for this – we have a primary site in Toronto and in the Toronto site we have two – we have a high availability infrastructure, that means we have the registry and the DNS infrastructure and that's mirrored inside that site. And then in Ottawa we have a backup site and in there we have a similar architecture.

So what it means is Toronto, potentially, we have two signers running in parallel and the backup site we have the same thing there. And then remote – in Ottawa corporate office – we have another site and that's where we're going to assign the keys. The challenge is to get the staff internally to understand the impact of doing security management – all that key stuff – to make sure that we don't screw it up. It's got to be done right; it's got to be supported right; its' got to be managed, monitored, and all of that. So it looks simple but it's fairly complex, from our point of view, to implement.

We have a preliminary schedule – so we started the planning – the dates you have we're planning to be fully deployed by February of next year sometime – 2012. So we have the high level bullets of the plan that we want to do. So now what we're doing is we're building a detailed project plan to get there and then build a lab internal, do some training, tons of work to get there. So it's a fairly large project to get executed.

The key stuff – so we got our tentative parameters for DNSSEC. So, that's kind of the somewhat a best practice that we want to implement. And we'll need to design and develop the architecture around those parameters. Now I'd say within the IT team at CIRA, half the team is software development job – so we build our own stuff. The thing here is that the DNSSECs implementation, we want to put that as part of our Agile process and as part of our Quality Assurance process. So we need to make sure that the code goes into production, actually goes through our lab performance testing, stress testing and all of that.

So it's a fairly huge – well it's not fairly – it's a huge investment on our part to build the lab infrastructure to simulate the DNSSEC environment. And that's where we're going to get to learn internally how to deploy DNSSEC within the lab environment; gain the experience; gain the expertise; build processes; understand how to deploy and manage all of that – and then we're going to start thinking about putting it in production. So there's a lot of work. It's not as easy as it looks to implement for a ccTLD. If you want to put all the process and methodology around doing it right.

One of our key objectives for CIRA, for .ca, is operational excellence. And a lot of the stuff we do is around operational excellence and that implies we can't just wing it, we need to do it with the right process. Risk management – as far as I can tell we're very immature technology right now that we're implementing. There's a lot of bugs. There's a lot of issues. So we need to be careful internally on how we design our solution, how we pick the architecture, which module we use, what kind of HSM technology we're going to implement, how do we do a high availability infrastructure, how do we build a resilient infrastructure within the site.

There are things that are not a lot of DNS experience out there. Within this room there is, but outside in Ottawa none of the IT system integrator, none of the IT shops knows about, they don't even know what DNSSEC is basically to start with. So we need to reach out to get knowledge and it's not easy. From a risk management – the risk register – we get an email once in a while with service impacting outages; I've seen a lot of those. Having a service impact outage doesn't jive with our operational excellence thing. So this only reinforces our fact that we need to do this right, with the right process, the right methodology and get it implemented right.

So, we're committed to doing it. We're going to do it right. It's going to take whatever time it takes, but we'll get there in a timely fashion and controlled. And that's supposed to be a little happy face.

Steve Crocker: Great thank you. Matt?

Matt Larson: Matt Larson from VeriSign. So I'm the last presentation before lunch, outstanding. Okay, so this is an update on what VeriSign has been doing DNSSEC wise, particularly around .com and .net. So let me give you a timeline – I'm including .edu because VeriSign operates the .edu, the technical backend of that registry under contract with Educause, which in turn has the contract to run .edu itself. And that zone is held in the same registry system that also hold .com and .net – that's for historical purposes and we've never seen any need to remove it so the advantage for .edu is that they get DNSSEC support along with everyone else and they also get all the .com and .net performance and redundancy. So they were willing to go first and .edu, as you may know, was signed in July and the DN record went in the root, as you can see, on July 29th last year.

.net was signed late last year – the DS record was published on December 9th and .com is signed right now – I'll just pause and let that sink in. But it's unvalidatable and everyone is probably tired of hearing about this unvalidatable technique, but I will describe it yet again in a moment. But the good news is we are on target for getting the DS for .com in the root on March 31st, barring natural disasters or unforeseen circumstances, we intend for it to be there.

We had hundreds of people do that work and I get the applause – outstanding! Thank you.

So, that was the update – what I want to do for the rest of the presentation is talk about some of the challenges that went into doing this and talk about the design because I'm hoping that that's a topic of interest to people. So the challenges are I think obvious – we need to sign and maintain a zone that's being continually updated; lots and lots of churn every day 24-hours a day to .com and .net. And we have extremely tight SLAs, for example, the SLA to create a zone in EPP is 50 milliseconds, this is actually the .net SLA. The .com SLA is considerably longer, but it's the same registry system so we hue to the same SLA for .com as well as .net. So we've got 50 milliseconds from the time a registrar initiates the create command to the time that we need to complete that create and acknowledge it.

And we have three minutes to get a change made to the registry from our data center out for all of our resolution sites or name servers worldwide. So that's a challenge. And obviously we need to safeguard the cryptographic materials. There's also a DNSSEC impact on resolution – obviously performance is going to take more bandwidth. I don't have the specific slide about that but what we're seeing with .net is roughly twice the bandwidth usage.

Then there is also a network issue regarding fragmentation in our environment; and I'll talk about that in a moment. And then one challenge that we took particularly seriously is that we've got to ensure valid DNSSEC responses. We just – we've built a lot of safeguards into the system and we don't ever want there to be a bad .com or .net or .edu DNSSEC response. So there's a lot of validation built in.

So, I guess that's reasonable legible, but this is a diagram of just the provisioning side. So this is just the registration side of the registry. So in the upper left there are, what we call, application servers – so these are the EPP servers that the registrars connect to. Everything new for DNSSEC is in a colored box. So it you'd ignore the colored boxes just for a moment, the EPP servers speak SQL to the registry database where things get updated. There is a continuous process that we call extraction where the incremental updates are prepared – they're validated after the extraction. And this is a validation, I should point out, this is pre-DNSSEC. This is a validation to make sure that what is in that update for example when it's applied to a local captive instance of what's out in the field, matches what's in the database and doesn't cause that captive name server to tip over or anything. So, only after its passed validation does it get pushed out to our resolution site.

So then enter DNSSEC. We went back and forth on the architecture we wanted to use. One possibility that we considered first was as EP transactions came in take the ones – see not everything is affected DNSSEC wise. If someone is changing a domain that doesn't have a DS record, then there's no DNSSEC impact there because of opt out. So one possibility we considered is having a separate queue of pending changes that were affected by DNSSEC. And then having something come into that queue, pull things out, sign them and stuff them into the database. So, sort of having two streams going into the database.

And after really looking at that for a long time, we decided there were just too many corner cases and we were concerned about just how complicated that would be. And the reason we had been looking at that solution was because we had assumed that being able to do all of the signing necessary in that tight SLA, in the 50 millisecond window, wouldn't be possible and would be very difficult. But we found out, we found the right HSM vendor and with after a lot of engineering were able to do the signing necessary within that 50 millisecond window while the transaction was going on. So for example, if someone adds a DS record – in 50 milliseconds we'll accept the DS record, we'll sign it, we'll commit it to the database and then we'll acknowledge the transaction.

So if we look at the pieces then that are colored, all of the signing – if you look at the yellow box – all of the signing happens through a signing service that abstracts the HSMs and more on that in a moment. Everything in green, are signing processes that aren't in that 50 millisecond path. So if you see from the application servers there's an arrow to and from the signing services in the upper left – that is the path that happens constantly as activity comes in from registrars.

But there are other DNSSEC signing operations that have to happen periodically. One – if we go down the list in that green box – one is the resigning process which is simply refreshing signatures before they expire. Another is the very periodic rollover process – the idea was to automate all of this so that none of this needed human intervention; at least not to kick off the process.

The SOA is special because we do bump the SOA for every one of our changes and we do create a new batch of incremental changes every 15 seconds. And they typically make it out to everywhere that we have an authoritative server within 45 seconds, usually much less. So every time the SOA gets incremented it has to be resigned. And then there's what we call a key signing request generator. So, the DNS key set itself is signed, but with the key signing key. So everything in the yellow box – the signing service – those are the HSMs that have the zone signing key; and they by necessity need to be online. Of course we've taken a great deal of precautions; they're not just under somebody's desk. I hope. Maybe I could rewind and not say that. But no, in all seriousness obviously there's a great deal of care surrounding that yellow box.

But the key signing leys do not need to be online; there's just no reason. We're rolling the zone signing keys every three months so the key signing keys only need to be taken out once a quarter. So that's done in an offline fashion and that's done reusing the scheme that we developed – that ICANN and VeriSign developed for the root zone. So I'll cover that in a moment.

So on the provisioning side things that had to change – I've basically covered this already. I didn't specifically mention that we are using RFC 5910 pretty much unmodified the EPP extensions for DNSSEC to pass DS records. We are also only excepting DS records – not DNS keys. The registrar has to calculate the DS record and send us that.

Let me talk a little bit more about the signing service and the key management. So we need multiple HSMs both for performance and redundancy and therefore we abstracted and we put the signing service in front of them and everything that needs to be signed interacts with the signing service and only the signing service knows about the HSMs. So that required custom software and we put high availability features into that. A little bit more than about the KSK management – we have a separate group called our cryptographic business operations group and they're the people who have the key ceremony room and who know how to handle sensitive crypto material and keep it safe. Speaking of safe, they have safes; that's those folks. So they handle all the key material.

As I said this reuses a lot of what we developed for the root signing project. So when the zone sending keys need to be signed they're packaged up in an XML document called a key signing request, similar to a certificate signing request in concept in the CA world. So that key signing request is sent to the CBO, they validate it, they sign it and then they send back the signed DNS keys in the signed key response – the SKR. So all that can happen in a somewhat manual fashion because – in fact, it has to because the HSMs for the KSKs are in a locked room that only the CBO can get into; so they have to put the key signing request on a USB token and carry it in and carry the SKR out.

So I've – let me skip straight to the slide about the signing server architecture. We have a single signing server that talks to a single HSM and then anything that needs to do signing talks to one of the signing servers and in fact, they know about all the signing servers, but because of this HA protocol a signing server can drop out and there's still another one to handle a signing request. And this is what I basically just said about the offline keys. One thing that is interesting is the cryptographic business office manages not only the key signing keys but also the zone signing keys. They generated the zone signing keys, they configured the HSMs that we have online and they hand carried them and tracked the chain of custody to the data centers where those need to reside. So if it involves keys, we're having our folks who know about handling keys and specialize in that, we're having them handle it.

Here are some of the parameters similar to what I've seen other people using – a 2K bit KSK – our intended lifetime is years. We don't have any specific schedule right now to roll it. So we'll roll it when circumstances demand it and we think it’s necessary. So we’re not committing to any particular timeframe when we roll it. We’re specifically not going to use RFC 5011 – the intention is if you want to trust the .com key, you should go to the root. We spent all this time and effort on the signed root, so trust the signed root, trust the DS record for com signed in the signed root and that will tell you how you can trust the .com KSK.

1024 bit ZSK rolled every three months; seven day signature durations – the refresh time in overlap is slightly different whether it was made with the KSK or the ZSK, but basically we’re signing at about the halfway mark. And RSA/SHA-256 and then of course opt-out and that’s for the reduced zone size not the confidentiality because anybody can get the .com zone if he wanted.

Here’s a little bit more – this shows the architecture now mostly – well I guess this is the bigger picture of the architecture that includes more detail on the resolution side. The components in blue existed prior to DNSSEC but got changes and then the single green component is new. So the extraction component that assembles the incremental DNS change and the validation component – those had to be made DNSSEC aware, but we added a specific DNSSEC aware validation function to that process.

And then ATLAS is the high performance name server that we developed for .com and .net – that’s been in production about 10 years. And that’s a two-layer architecture with a front end that does DNS protocol work and a back end that has a database with zone information. And actually the front end really needed no changes for DNSSEC; it was only the back end had to be made DNSSEC aware.

So the DNSSEC validation – this is that green box on the previous slide. We do the obvious things in that green box. We're verifying every signature – before we publish the signature we verify that it verifies. We check the sanity of NSEC3s – they obvious things that you would do if you wanted to make sure things were okay from a DNSSEC perspective.

Just a very little bit about fragmentation – because of the design of our networking equipment, fragments are a challenge. So we decided rather than dealing with a lot of network redesign, which would be yet another thing we would have to change and we're very concerned about stability here, we decided to create a solution where we just didn't go over the Ethernet MTU. So the design is we don't ever build a DNS response that would fragment if the MTU were larger than Ethernet ones. So we just stopped putting stuff in the packet and if necessary stuff the TC bit. We find in practice that just doesn't happen, the TC bit being set that is. We see a very nominal increase in TCP in .net.

So, everyone has probably seen this sort of slide multiple times. We are using the deliberately unvalidatable approach. We have this very cautious and deliberate approach. As you can imagine, nobody working on .com wants there to be any problems with .com from the DNSSEC perspective. And so we're using the same technique of this unvalidatable key. So as you can see, that's the actual key that's being published right now in the .com zone and it's deliberately obscured key material.

And what that allows us to do is an incremental deployment. The whole reason for that is so that we can do a very slow rollout – and here's how it worked from a resolution side. We had the new code, the new resolution code that supported DNSSEC, so we rolled out just the code very slowly without enabling DNSSEC. Then when we were happy that the code was stable, then we started turning on DNSSEC, but we did it one site at a time with a lot of baking and letting things stabilize.

That's where we are now. We have the unvalidatable zone being served everywhere. And the last steps, which will happen in the end of March – then are unblinding – we call it – the key. We'll publish the actual key. And then after a short interval we'll add the DS records to the zone. If we look on the provisioning side from a registrar's perspective, we always have an operational test and evaluation environment; that's always available to registrars; it's a sandbox where they can test their EPP code against a live server. And know that whatever is there will be either the current, or sometimes the next version of the code we're going to deploy. So we've had OT&E available with the EPP DS record, the DNSSEC extensions; that's been available for some time.

And in fact those DNSSEC extensions are available now in the live interface – a registrar could, if they wanted, submit a DS record for .com right now and we would publish it. Well the zone – we're no longer publishing in an unsigned zone; it is a signed zone. The question was would we publish it in the unsigned zone and there is no unsigned .com anymore.

So, in terms of – I'm winding up here – in terms of issues during deployment. The edu zone deployment back in July was absolutely just no issues that anyone reported to us. We had what I would call a very minor hiccup in the .net zone deployment and it was related to an issue that was uncovered with some versions of bind. And we've communicated that to ISC; ISC is on top of it. I believe it's patched in subsequent versions now. And this was a low impact, we only heard about it from one person, but when somebody said you signed .net and my name server stopped working we obviously wanted to chase that down and see exactly what was going on.

Lessons learned – I think the biggest one, and this is what I would list as the biggest lesson learned for the root DNSSEC deployment as well as what we've seen in .edu and .net, is that the internet didn't break. We've proven we can overlay DNSSEC on top of this and it works. We're very pleased with this incremental deployment technique that we've developed. This registrar test environment that we have – and we added the ability, we're publishing a singed zone from the test environment. So a registrar can see changes move end to end. The registrars know the IP addresses of this sign zone being generated from the test environment and they can fire queries at it and see.

I mentioned the minor issues we had with some installed base of hardware and software. So then, I would say what our best practices are would be this slow rollout, very strict key management practices that I've described with our cryptographic business operations unit handling them. We've decided that his online zone signing key and this offline KSK is the best model for us.

We do publish a DNSSEC practice statement – if you have some time on your ands you can read all the policies and procedures that we have regarding DNSSEC. There's a separate one for each TLD; so .net and .edu is already published and .com will be published before we sign the zone. And then we believe we have to do a DNSSEC validation of the zone data before we publish it.

One last slide which is just we've done a lot of work with the registrars. There's an SDK; we've always published an EPP SDK and we've added the DNSSEC support to hopefully make it easier for the registrars to add DNSSEC support on their end. I've already mentioned the OT&E environment. We've got a resources center that's available to registrars with a lot of information and

I’ll talk later today about the signing service. And with that, thank you.

Steve Crocker: Thanks very much. I want to – so this is time for questions and come on up to the mic, but I want to start off with a question for both of you and pick up from those last bits that you covered Matt, but I’m really addressing both Jacques and Matt here. Please say more about the engagement of the registrars and about what you see from the registry side, what the uptake process is likely to be in terms of the registrants signing their zones and the registrars – I guess I really might want to focus more on the registrars – the readiness of the registrars to accept key material from registrants and/or signing zones on behalf of registrants – either order.

Jacques Latour: As part of our plan is to have a certificate accreditation process for DNSSEC for the registrar. Today we had very, very few requests for DNSSEC, in Canada, I guess. And the state of our registrar with respect to DNSSEC is it’s not documented, we’re not sure.

Steve Crocker: How many registrars does CIRA have?

Jacques Latour: We have about 150 some around there.

Steve Crocker: And I would imagine, I haven’t looked at this, but I would imagine that similar situations in many places where there are a few big registrars and then a large number of relatively small registrars.

Are you getting any interest or interaction from either end of that spectrum?

Jacques Latour: No. We had about a handful of requests from registrants and that's about it.

Steve Crocker: And recalling what you were saying about your schedule, you're plan is to be in operation early next year right? So, maybe the interesting time to revisit the question with the registrars would be a bit closer to that point in time.

Jacques Latour: One of the – internally, the thing I'm trying to do within CIRA is the objective of the DNSSEC project is not to sign the zone. The objective of the DNSSEC project is to have a certain percentage of .ca sites that are important, like the banks and the target audience that would deserve that; that would be the success criteria of DNSSEC. That involves we need to go out with communication; talk a lot of media awareness; we have a Canadian Internet Governance Forum; we got a lot of stuff. So once we start the project we've got to do outreach and explain the benefit of DNSSSEC. So that's our strategy.

Steve Crocker: Good. Thank you. So Matt, no long delay between now and when you're going live so the question about the registrars is quite relevant right now, right?

Matt Larson: Yes. Well, first I should say in contrast to what Jacques said, we don't have an accreditation program specifically for DNSSEC; there's an overall accreditation program and a registrar has to pass a certain number of tests in our OT&E environment to get accredited in April and to be authorized, I should say is a better word, to do changes. But not specifically for DNSSEC, so therefore we don't know the exact number of DNSSEC capable registrars. I do know that it's not zero and that it's not just some of the smaller registrars. There are some of the very largest registrars that we know have expressed interest in DNSSEC and our implementing it.

And if I could just go back to my last slide, you know we have, the registrars have been a significant focus of our overall deployment because without registrar adoption and interest in DNSSEC to allow registrants to DNSSEC, all the effort on our end will be for naught. So we've tried very hard to engage the registrars and make it just as easy as possible for them to implement DNSSEC.

Steve Crocker: Would you want to make any kind of prognostication about either number of registrants or number of registrars that are likely to be on board with DNSSEC in some period of time; next few months?

Matt Larson: I really don't know. I mean that's – we've planned from a capacity standpoint – they can sign them all, they can sign all 100 million names, they can send us DS records for 100 million and we'll do it, but I don't know. Maybe some of my colleagues who are on the business side and much closer to the registrars would know, but I don’t want to even hazard a guess.

Steve Crocker: Okay, time to open the floor here. Rick.

Rick Lamb: Rick Lamb just as .com registrant. My question was somewhat technical – the validation you do when someone registers a domain under .com for example; can they submit to DS records – one that validates and one does not. And I ask this because when you do key rolls, and I’d want to do that with my little .com domain, I’d have the same issues that you guys would have about packet sizes. And if I could do two DS records in there it allows me to create a much smaller key set when I do a roll. I mean I could be wrong, if there’s some other way to do that I’d love to see it, but I’ve racked my brain a few times thinking about this and trying to figure out how I could do it. So the question is, can I put in two DS records – one that is valid, one that I’m going to – to facilitate a roll and will com publish it?

Matt Larson: Yes absolutely. SO to be clear, the validation we’re doing is our own signature, we’re not doing any validation of the data in the DS records passed to us. However, we are validating the algorithm field in the DS records. So if it’s in the IANA list of assigned DNSSEC algorithms, if it’s a legal algorithm in the DS record, we’ll take it and we’ll publish it. I don’t know what the number of DS records is, but it’s over 10, it’s a very large number.

Steve Crocker: No other questions? This is – uh oh.

Gavin Brown: Hi, I'm Gavin Brown from CentraNic. I had a question about CIRA and having an accreditation process for registrars to use DNSSEC – I'm kind of puzzled because obviously an accreditation process is designed to stop people from screwing up. But it occurs to me that the failure rate for a registrar screwing up one of their registrant DS records is the same as if they screw up their name server records. So I'm puzzled as to why you have that separate accreditation process.

Having talked to registrars who are interested in supporting DNSSEC, they find it hard to integrate it into their product offering because some registries are saying you can just do whatever you like, like VeriSign is saying, you can instantly have access to this functionality - some of the registries saying you aren't, and they want to be able to support the same functionalities for all of the domains that they support. So, having some registries that don't have a barrier to the use of DNSSEC functionality in the registries is slowing down the deployment of DNSSEC in those registries that don't.

Jacques Latour: I guess – first of all, we're still in the planning stage for doing DNSSEC, but today we don't know how many registrars are DNSSEC. So it's not an entirely new process for accreditation. The idea is that this is the list of registrars that are accredited, this list supports DNSSEC, this one does Ipv6, so it's visibility on…

Gavin Brown: So, it's not a kind of technical evaluation process where they have to prove that they now have to send DNSSEC records correctly?

Jacques Latour: Yeah, if you do DNSSEC, you have to prove you do the record properly.

Gavin Brown: Do you require that they prove that they can send name server records correctly?

Jacques Latour: No that they're able to – not correctly but that they're able to submit the keys. I'm not sure exactly how it's going…

Gavin Brown: My argument is that if you just turned it on for everyone, the ones that were able to use it would use it; and the ones that weren't able to use it wouldn't accidentally use it or use it wrongly.

Jacques Latour: I'm not sure – can you talk closer to the mic?

Gavin Brown: Sorry. The idea of doing a technical evaluation of a registrar before they can use the DNSSEC features of your registry seems to me that it's kind of keeping competent people competent in that if a registrar is incapable of correctly using those features, they're probably not going to try and use them in the first place.

Jacques Latour: Well, from our point of view we need to understand who does DNSSEC…

Gavin Brown: I completely understand that there's a communication thing and you want to be able to tell your end users these are the registrars you need to go to if you want to use DNSSEC on your .ca domain name – I definitely agree with that, that's a very good useful thing. But my question is about having a technical evaluation.

Jacques Latour: I don't know what the – today we're just planning like I said, so I don't know the extent of the planning is going to be for that.

Roy Adams: My name is Roy Adams, I work for Nominet. I would like to respond to the person who just asked this question. We had an internal discussion at Nominet if we should or should not validate DS records when we get them from our registrars. We eventually adopted the concept of garbage in/garbage out. We do check, just like VeriSign incidentally, we do check the algorithm field, but that's it. If people send us broken records then they send us broken records. The problem is if we go and check them, five minutes later they might be wrong. If we go and check them and then tell the registrar "no we can't accept this" – we just are another bump, another hurdle for registrars to take. I'm not sure if I'm answering the question, but I'm just another hurdle for registrars to take.

Steve Crocker: Sorry, I think we've discussed two issues in the last few minutes – one is does the contents of the DS records themselves get validated to make sure it's a DS record corresponding to a key published in the child. And the other is, are registrars themselves validated to ensure that they can properly do EPP with DNSSEC extensions.

Roy Adams: Oh, this is typical Roy misunderstanding the question. I'm sorry.

Steve Crocker: But I'm glad to know what Nominet does.

Gavin Brown: I was going to say, Roy that I think you were agreeing with me.

Jim Galvin: Jim Galvin from Afilias and I guess I want to partly respond to the question too if I may – I mean we're the back end for 15 TLDs and .org in particular that offers signed delegations. So we do require DNSSEC OT&E; the registrars to validate themselves. And the model from our perspective is it's not about the individual data so as Matt, correctly distinguished between the data in the records and the OT&E in the large, we don't validate the DS record against the key or anything like that.

Registrars are responsible for the data, they have to do the right thing or not. It's more about recognizing that the registrar properly handles the EPP extensions with DNSSEC data. So do they handle all the error conditions and all of those tests and do they respond correctly to them. So can they submit valid transaction and then can they also handle all the error conditions that go with it. And that's really the basis for our testing is just demonstrating that they have that knowledge and expertise. And it's just consistent with our overall principle of very structured, very careful deployment.

Steve Crocker: Thanks. Rick.

Rick Wilhelm: Okay, Rick Wilhelm; Network Solutions. Being a registrar, one of the things that we find is that it's useful to decouple and let us go through technical OT&E and separate that from being able to push records in using a web tool or something like that. Because what we find is that to require the technical integration adds to the cost and puts another barrier because, like right now, we have a few customers that are coming to us and they want to do DNSSEC and so we handle it over the phone because that's, right now that's the cheapest way for us to handle it. Long term – that's not the right solution, but it's a rolling start.

Matt Larson: Right. So today you can use the VeriSign registrar tool to add records and you don't need to use EPP.

Rick Wilhelm: Which we do and we do so joyfully and we appreciate that.

Steve Crocker: Thank you. Andrew.

Andrew Sullivan: Andrew Sullivan. If I out my IETF hat on for a moment, I will point out that there is a technical reason why registries might want to be doing things to the DS that they're not doing to other kinds of data. Because the DS record is authoritative data from the parent and not from the child; whereas the NS stuff is stuff that you're on both sides of the zone cut. So there really is a technical different here about this data, and so there may be in fact a reason why the parent may want to validate that data.

Steve Crocker: This is absolutely perfect. The questions have subsided at the exact moment that we're supposed to go in and have lunch. In principle, everybody is supposed to have a ticket which you should have received coming in. Julie, you want to say a bit more?

Julie Hedlund: Yeah, I realize that some of you may have come later and not gotten a ticket and some people have left who perhaps had tickets – the idea is that we don't want people coming up the elevator to partake of our free lunch which you've all worked so hard to have. So since you've made it through the workshop so far, thank you very much. Please go in, I'm going to go in and work with the person at the door to make sure that everybody that's here gets in for lunch. We can accommodate you. Thank you.

Steve Crocker: What's the status about security in here during lunch?

Julie Hedlund: There is no security in here during lunch. I suggest you take your various apparatus with you.

Steve Crocker: Lunch is on the other side of the hall here – same floor, other room.

Julie Hedlund: Yeah, it's Victor's Room – you can't miss it just walk past the stairway and there should be somebody there asking for your tickets.

Steve Crocker: And come on back here at 12:30 – we're going to start up.

[break]

Julie Hedlund: Excuse me momentarily. I would like to ask the panelists, if you are in this final panel to please come up to the front table and take a seat, and we will be starting momentarily. So any of the panelists who are here please do come up to the front as quickly as you can and we'll be starting momentarily. Thanks, everyone.

Thank you, everyone, for coming back after lunch for our exciting final panel. And so I would ask all of you to please take a seat and we are going to get started here in just a moment, so please come in and take your seat and thank you for coming back to join us.

Male: While everybody's coming in to take their seats I think we should give our lunch sponsors one more round of applause.

Steve Crocker: So this is the last session and it's focused on signing services, and as I mentioned at the beginning we're going to change the order from what was published slightly. Matt Larson has to be at two places at nearly the same time so we'll start with Matt and then we'll proceed with the original order. Matt, you ready? Good, thanks.

Matt Larson: Okay, thank you Steve. I wanted to talk about VeriSign's DNSSEC signing service, so next slide, please. You're probably going to hear over and over again what a signing service is but not surprisingly, it converts unsigned zones to signed zones. Ours, like probably many others, is a bump in the wire using zone transfer, so we pull the zone from the customer's designated server, sign it; they pull it from us and they do whatever they want with it. So we only sign; we're not hosting it for this particular service.

There's nothing to be done after setup, and the manual setup, the provisioning process is pretty straightforward. So in terms of the bullets here, what it provides: obviously signs the entire zone, key rollover is automatic. I really get the sense you're going to hear this same presentation several times in a row – I'm just, I'm the lucky one. I get to go first. I'm getting the "Matt, can you just say what's different in yours from everybody else's?"

We're committing to three nines uptime. We're also committing to an SLA for a ten-minute signing, so in other words from the time that the customer sends us a notify we'll send to them a notify back within ten minutes, giving them their sign zone back. But in production as you can see you know, we're running 60 to 90 seconds for that so much faster. And everybody has their own keys; we're using good crypto hygiene here, we're not sharing keys.

Here's the architecture, nothing remarkable here. There's basically two management lines here – one to the unsigned master… I mean the reason this says "registrar" is that it, I'm getting a little bit ahead of myself with this slide. The service presumes that this is going to be operated by registrars. That might be what's different than everyone else's – the target audience, this is for registrars as I'll explain for a moment.

So the reason there are two gray management arrows, the one from the registrar who needs to gain some weight pointing at the master server: that is to get the zone content there and then the one pointing to the red cloud is to actually provision the signing service, saying "Please start signing acme.com and pull it from this particular server." Next slide, please.

Oh, alright. So those gray management arrows, there's both a SOAP web service or a web-based Gooey so the registrar can take their choice. In terms of some DNSSEC details, we're resigning the entire zone every single time so the presumption is that most of these zones are going to be small enough that it's just easier to resign the entire zone than it is to do some sort of incremental thing. And every time we resign it the signature validity period is 14 days, and then we do refresh at seven days. And I don't need to read the details to you on the bottom of the slide – it's very typical parameters for key size and rollover intervals.

We do give the customers a choice of NSEC or NSEC3, and we are storing the keys in [fbps 140/2-level 3 HSM]. So that's a decision we decided that even though we're going to support lots of (inaudible) we wanted to use HSM to do the right thing for key storage. Next slide, please.

So it's available to ICANN-accredited registrars. The goal of the service and I suppose I should have said this on the first slide rather than the last slide, is to make it easier for registrars, excuse me, to make it easier for registrars to DNSSEC-enable their hosting services. So the presumption is that there might be registrars out there who already have a DNS hosting service that they operate but making that DNSSEC-capable could require a lot of work. So instead, if they're willing to put a bump in the wire zone transfer base solution in place then the idea is with fewer changes they could DNSSEC-enable their hosting service.

So this is all part of our attempt to make it as easy as possible and to be as encouraging as possible to registrars when it comes to DNSSEC adoption. You know, we didn't want anyone to be able to say "Well, I would add it, support, but then I have all these other customers, I'm also hosting their zones and if they DNSSEC enable then I can't host their signed zone."

So there's a free trial period, absolutely free as in beer and speech, and that's for, those are for VeriSign TLDs; and if the registrar wants to take… If they have .org for example, if they're hosting a .org domain that will cost them $2 per year, and that's it. Thank you.

Steve Crocker: Excellent. So in deference to your schedule here, let me ask if there are any questions for Matt and then we'll hold off questions for other people until we're at the end. And so come on up; we'll do… Do you have time for-

Matt Larson: Absolutely, thank you.

Steve Crocker: Okay, thanks.

Matt Larson: Thank you for accommodating me. I appreciate it.

Roy Arends: Is this on? Yeah. My name is Roy Arendsen and I'm from Nominet. Matt, one question: are you going to store the actual private DNS keys on the HSM or are you going to use an encrypted store where you put the key of the encrypted store on the HSM?

Matt Larson: No, the keys are actually stored in the HSM as… No, we're storing the encrypted blob of the key outside and it's send to the HSM for the actual signing.

Roy Arends: Okay, got you. I mean that's, that to me is secure enough but this is the way you gain speed.

Matt Larson: Right, we're storing these opaque, the keys become encrypted opaque blobs.

Roy Arends: Okay, got you. Thank you.

Sabine Dolderer: I have a question. Do you limit the (inaudible) file size of the domains?

Steve Crocker: Can you state your name so we can put it-

Sabine Dolderer: Okay, Sabine from DENIC. Do you limit the sole file size of the domains you sign and does it affect the SLA you actually propose?

Matt Larson: I don't know the answer to that. I don't know if any of my colleagues… I think I'm here solo. I think the two people who could have answered that question are now not here so I apologize. I would imagine there's a limit but I don't know what it is.

Steve Crocker: Good. I think we've reached a stable point here. Thank you very much.

Matt Larson: Thank you.

Steve Crocker: And this will of course make a big difference, the signing of .com and unveiling of all of this. So when we reconvene in Singapore there's going to be an ability to look back just briefly and see where we are in all of this. So thank you.

Before we proceed with the rest of the speakers I want to put in a little commercial. The, a lot of work goes into planning these sessions. We have had the benefit of two volunteers over the past couple of years who have sat through the sessions and come up to me afterwards and said "Eh, this was okay but we could make it better." I said "Well great, you're elected here." So Marcus was first and then Simon fell into the same trap and we swallowed him up as well.

Marcus has wandered off and we have openings, so any of you who are, who think you've got useful ideas about how to keep these sessions live and fresh, please don't be at all bashful or hesitant. And it's a good group – we enjoy working with each other, and the workload is approximately a call once a week and sort of incremental planning. And when the meetings are bunched up and coming a little faster than usual, like the Singapore meeting is coming in three months, we tend to look ahead a little bit so we're doing a little bit of advanced planning.

So feel free to raise your hand or come up quietly, or just be a little slow about leaving the room – we'll catch up with you. Thank you. So moving right along, Jim Galvin from Afilias is next.

Jim Galvin: Thank you, Steve. Okay, good slides. Next slide please. So, last year in June when we had launched signed delegations in .org, obviously the next big step was to think about what it would take to get registrars on board. Obviously we had made good progress even since then because the root has been signed; now the number's actually over 70 TLDs are signed, and as Steve was reporting earlier this morning there's variability as to what "operational" really means, and hopefully we'll have some more consistent data going forward with those statistics.

And there's still a lot to come, and we were focused specifically on registrars because obviously a lot of the other things don't make as much sense if you don't have signed domains themselves to work with. So next slide please.

We, Afilias actually conducted a survey of registrars in 2010; our goal was just trying to understand that channel and what those people thought about DNSSEC. Some of our findings are generally people think it's a good idea, but preaching to the choir here; and a lot of people will say something's a good idea because you can always make it sound good when you're trying to sell it to them. The question is whether you can actually get them to want to do anything with it.

It's clear that people…There's a fair amount of expertise that still has to grow and be available, and we were able to identify some issues in the registrar channel as to what it means to deploy DNSSEC. Now, I apologize – with these slides unfortunately one detail I forgot to include on the slides was an actual pointer to this readiness report that we had done. It is actually free, freely available. If you go to Afilias' website, www.afilias.info, and you go to the search box and you just say "registrar DNSSEC readiness" it'll take you right to the report, and you'll be able to find a link where you can actually download the entire PDF and see it. But next slide please.

I pulled out a couple of the statistics. It shows you all the questions in this document and all the actual results that we got, but I wanted to follow a couple of things here which I think are still relevant. This was still only six months ago. It's interesting here, the question of asking registrars "When will you offer DNSSEC?" and you find that 32% have no plans to offer it in the next twelve months, meaning basically all of 2011. Slightly more, 37% say that they will have something in 2011. Next slide.

But what's interesting is in spite of the fact that they think they want to have something they really don't know what that is and what they're going to do to get it, so we had two particular questions here about, you know, "Will you buy it from a signing service or will you buy it from a registry or DNS provider?" And it's pretty telling in my mind the percentages there – 45% don't know whether they'll buy it from a managed service provider, and you've got 62% there that don't know if you'll buy it from a DNS provider or a registry. So I think that people have a desire and an interest in wanting to do the right thing but they don't know how to do that. Next slide.

So I guess what you're hearing up here is we all have our own version of a signing service to offer, and we're calling ours One Click DNSSEC. So next slide. The feature about it is in fact that it works with just one click. If you get the managed services from us that have DNSSEC included you can just have your domain name and you click in one spot and everything is taken care of – all your key management is done, your key rollover, coordination with registrars.

I mean since we're a registry service provider offering that service then with the right relationships we can actually coordinate the insertion of your DS records in for you and make all of that work with the system, too. So that becomes a nice way to fill that gap - as others this morning have identified there is that gap in the relationship between your DNS provider and your registrar in getting your key material up into the registry. Okay, so next slide.

This is just a quick look for completeness in the set of slides about who Afilias is. These are the 15 TLDs we support and 18 million domains under management. And then the last slide if you move on, if anybody has any questions or issues that we don't cover here we do have a booth downstairs where other folks who are marketing people – not me – will be there and answer your questions. So if you have a technical question you should bring it up here and I'll answer it for you. Thank you.

Steve Crocker: Thank you. So moving along quickly, Peter from EURID. Thanks.

[background conversation]

Peter Janssen: Good afternoon, just waiting here a second for my presentation to be found again.

[background conversation]

Peter Janssen: Okay, good afternoon. For those that don't know EURID it's the European registry, so just one of the other registries that's sitting on this table today. Actually all of my presentation is sort of the same thing as has been said before and probably will be said after me, so I'll try to concentrate on the little bit of differences to anticipate the question from Steve to actually highlight differences instead of the similarities.

Just to introduce where we're coming from and where we want to go, this is a typical setup. You have on the left side, that's your left side as well, yes, a registry in sort of orange; on the right side you have a registrar. They have a whole lot of infrastructure in place to actually register a domain name and actually make it work out there on the internet. So you have a registration engine that puts it on the database and then for .EU's specific case we have something called [Anamic Updater] that continuously watches the database and feeds it to a hidden master that then actually feeds it to public slaves, with or without the DNSSEC. We added the DNSSEC part so the Anamic

Updater actually feeds it to the hidden master that signs it on the fly and actually pushes it out.

On the other side you have the registrar that has a very similar process in place. It has a provisioning process that actually does the DNS nitty gritty details of getting some sort of zone in place that's pushed out to the public slaves, and on the front end you have something called a registration engine that actually talks to the registry site file EPP or a secure website – in our case to actually register the domain name.

So if a registrar wants to add DNSSEC the whole picture becomes more well, filled with all sorts of arrows and blocks and things like that. The (inaudible) of it is that actually there is something called a signing process that actually sits in between the provisioning posts as in the actual zone, getting out of the internet that actually takes care of everything; but the important part you have to notice is that operant red arrow that actually talks somehow to the registry to have the keys that actually are used to sign the zone actually get to the registry to be included as DS records in the parent I would say.

So moving along, the DNSSEC signing service as envisaged by EURID is that actually what we do or want to do is the bump in the wire as it's been called before. We break that link from that registrar hidden master to feed its public slaves and actually deviate the zone to a zone signer module that actually feeds it to a hidden master that then pushes it out again to the public slaves. So we take in the zone, sign it, and push it out again to the public slaves. So with a registrar it's just one small step, get the zone to the signing module of the signing service and actually receive the signs down again from the hidden master.

And also all of the extra arrows, most importantly the purple dash-dotted line from the zone signer to the registration engine of the registry because their signing service obviously has the same problem if it generates a key and finds a zone (inaudible) than if it's a case that they actually need to push to the registry for inclusion in the parent. So all that will be taken care of very much as Matt before me actually explained. The little blue arrow, the dashed line from the zone signer to the public slaves is just one extra check to make sure that actually the signed zone is out there before their keys are actually published in the parent zone.

So what are the goals of the DNSSEC singing service as put in place by EURID? The idea is to reduce the need for increased resources any which way – it could be human resources because a registrar has some sort of a learning curve to go through. He needs to get his mind around "What is DNSSEC? How do I cope with it? What do I do?" and so on and so on. There is the admin part of it, the continuous resigning zone key management, key rollover and all that sort of thing, and obviously there will be some sort of a hardware impact – you need some machines to actually do this.

Another goal is to minimize changes to the registrar infrastructures, both on a DNS platform as well as on a registration platform to actually make as little changes to anything a registrar has had in place for probably the last ten or fifteen years, make as little possible changes whatsoever. Automation is one of the key words to actually prevent human error: if you automate it you will have some sort of a guarantee it always functions in the way that you want it to function.

Obviously security is a big issue. If you're going to sign zones you actually want to make sure that the zone that gets fed to you to be signed that that actually is a real zone and not something else, and vice versa. So that's one of the key aspects there. And as long as you can keep it on a standard level – that's things that are actually used by other registries and registrars out there in the world – that would be great.

So what would the typical flow be? During domain name registration the registrar actually provides sort of a flag to say yes, this domain name should be under management of the signing service or should not, so the classical trade-off yes or no. Also there should be some sort of information about the hidden master, where our signing service and actually go and transfer in the zone file.

This is the case to actually make the zone transfer secure, and actually also the IP address or addresses of the public slaves where the signed zone eventually gets pushed to, and again, these are the keys to make that transfer secure and safe again towards the public slaves of the registrar. And from that moment onwards actually it's the signing service that takes control of the whole nitty gritty admin I would say.

So where are we? .EU got signed in September, 2010. At that moment we got the DS in the root zone; actually this zone was signed quite before that. We also did a sort of staged way for us to be signed to zone that didn't accept any DS records from registrars; then we added them and they got included in the root zone. We have sort of an improved concept phase where we're designing and building this system. It's not quite done yet but we are close. We're going to deploy it in the live .EU zone so we want to actually make this happen on the real .EU thing, and obviously to test something in real life we have a few registrars that are actually very interested to see this happen and actually want to work with us to see what and how and where.

The last thing I would say are some potential side results. There you saw that the registrar needs to mention the IP addresses of the hidden master, the public slaves, that gives the keys the little flag to split it on or off. Actually, one of the ways of doing this is to communicate via EPP so we envisage to have some sort of an extension that might be standardized in IETF, ways to actually you know, not reinvent the same thing each time when you progress down the table here. So it might be nice to have actually something that is the same in whatever registry that you talk to.

Another thing is the DNSSEC delegated key management interface. That's actually the interface that a signing service will use to talk to our own backends to actually upload the keys, manage the keys and things like that. As long as we're building that we actually might make that available to also, for instance, the DNS host which is not necessarily the registrar. So in that case the DNS host doesn't have to pass fire, the registrar actually can do the key management but can work directly, talk to the registry to do upload of DS records.

And a potential side result, a last one I would say and something that's rather hot in the community as well: "How do you do transfers from one registrar to another one when the zone is signed? How do you do that without losing, resolving and losing signing?" If we as a registry are the signing service provider for the domain name that actually gives us the opportunity to actually makes transfers ripple free in the DNSSEC signed zones.

And I think that was the last slide? Yes, thank you.

Steve Crocker: Excellent, thank you very much. Simon? The signing service for Nominet.

Simon McCalla: Thanks, Julie. Next slide please. I was rather mischievously trying to think what would differentiate our service from everybody else on the panel, so I think we're not as far advanced as everyone else, we're still in design phase, so I reckon it's going to be a very British DNSSEC service and our records are going to be signed by the Queen rather than by- (laughter) and you have free fish and chips with every iteration.

The reality of it is we're the same. It's a bump in the wire service. It's going to work in the same way as most other folks' are going to work. I think it's a good thing that these services are going to be pretty standard. So basically we're the same, we take unsigned zones and we return signed zones. Next slide please, Julie.

So I think just a question for us is why are we creating a signing service? Well, we wanted to firstly increase DNSSEC adoption across the UK. We've got a relatively skeptical registrar population about DNSSEC. We wanted to make it as easy as possible for them to adopt this and get moving, and they've been very positive about the noises we're making around it. And we see it benefiting for us both our large and our small registrars. I think our largest will use it as a kind of stepping stone service, so it will confidently allow them to dip their toe in the water with DNSSEC, try it out, see how it flies if you like without investing significantly in infrastructure.

And we think the small folk, bear in mind what's unique about the UK is we've got 3500 registrars, and some of them have only got 10, 100 domains. So the small folk will use this as a very cheap and straightforward way to implement. Next slide please, Julie.

Oh, that's strange – we're missing some graphics. That's a bit odd. Never mind. It's the same service really. Fundamentally we've got two flavors, basically. We've got what we call Simple Signing Service, and this is from a plea from our registrar community to make this as standard as possible and all them to possibly not just sign their .co.uk domains but also possibly other TLDs as well. So very simple we're fundamentally taking an unsigned zone and passing back a signed zone, and allowing them to put their DS record where they want to. Next slide please, Julie.

Again, strangely missing, never mind. The other thing is our sort of more complete signing service will also insert the DS records into the UK zone as well, so as you know we've got .uk and .co.uk, so that will provide the additional service. Next slide please.

One of the things I wanted just to address is obviously Nominet amongst a number of other folk here have been involved in designing the open DNSSEC product, and I think when we announced we were going to launch something like this people were saying "Oh, you've already committed to open DNSSEC. Does this mean you're not committed to open DNSSEC?" Absolutely not. We are fully supportive of both and we see them both having an equal place. We see open DNSSEC as a great solution for people who want to create their own DNS infrastructure and manage their own keys, and use that; and we see a signing service as a way of getting DNSSEC up and running quickly. And we foresee it being quite possible to start with our signing service and then move to an open DNSSEC infrastructure, so we're completely supportive of both ways. Next slide please, Julie.

And that's it for me, thank you.

Steve Crocker: This is moving along very, very nicely. Bill, you get to bring up the rear here.

Bill Woodcock: Okay, so again, a focus on the differences rather than the similarities. I think the main difference with us is the type of organization that we are. We are a global not-for-profit infrastructure support organization rather than a for-profit DNS services organization. So as a result that means that the service that we're providing is focused specifically at ccTLDs, not registrars; and we're very much aimed at doing knowledge transfer and helping the ccTLD gain operational experience and eventually transition into operating their DNSSEC signing internally rather than having an ongoing dependency on us.

So this was a result of a bit of a collaboration which took the form of ICANN allowing Rick Lamb to work with us to bring this service up, which allowed us to replicate exactly the root signing mechanisms and practices; and in turn we're taking the whole tree of documents… As most of you probably realize, the hard work here is in the key management and the business practices and so forth; the technological part of doing a DNSSEC signing is real work but it's relatively straightforward.

So we are taking that entire tree of documents and open source publishing it under [greater commons laws] so that anyone who wants to use that can satisfy their auditors without having to generate it all themselves and go through their legal reviews. Yeah?

Rick Lamb: Well, I'll just talk this way. Hi, I'm Rick Lamb from ICANN. I just wanted to emphasize how we arrived at this. One of ICANN's stated goals lately has been to accelerate DNSSEC deployment and we were looking for ways to do this while still maintaining a good system, a trustworthy system. And so this was a perfect opportunity for someone that had the right security models in place – namely PCH – to do this.

The relationship with ICANN is just like Bill said – it's just a little bit of my help, there's no formal contract or anything like that. His organization is just very well-placed to help this happen and it is very much of a stepping stone. I'll let you go back to that.

Bill Woodcock: So again our approach is to have this shared secure signing platform that any ccTLD can use that includes a knowledge transfer component to help them begin to take on each of the parts of the process themselves as they're comfortable doing so. The idea here was to replicate the root exactly, largely so that auditors would not find any separate questions to ask, right? If they're happy with the process that the root signing uses then they'll be happy with this one; no real differences.

Like the rest of our services this is provided without respect to ability to pay. We don't charge for these things so consequently it's funded out of our overhead and we hope that it doesn't grow too wildly successful.

A modular system, again so that people can take on holding the KSKs themselves, generating the ZSKs themselves, handing those off to us. Benefits I've already sort of gone through…. A transition path back and forth with documentation and checklists and so forth to make auditors happy about how it is that a ccTLD would move onto the platform and how they move off of the platform, and guaranteeing no lock in and so forth.

The platform probably looks like many of the others, again, particularly like the root one for level 4 HSMs: two online signers in Zurich and San

Jose; two offline KSKs in San Jose and Singapore. So we're trying to have diversity in locations. The Zurich location is being hosted by SWITCH, the Swiss Research and Education Network; the Singaporean location is being hosted by the Singaporean Communications Ministry; and the Americas location is being hosted by Equinix, so we have a governmental agency, an educational agency and a private sector. And obviously this is kind of building on top of our global Anycast platform, so there's a degree of integration if you want it; if you don't want it they're completely separable components.

So again, the level 4 HSMs sitting next to signers or the KSK masters inside a class 5 IPS, that in turn is inside a [skiff], that in turn is inside a tier 4 data center or the equivalent in the governmental or educational facility. The San Jose facility is there; the Zurich and the Singapore facilities should be built out by the time of the next ICANN meeting in Singapore. We're obviously targeting being able to do a key signing ceremony there at the next ICANN meeting.

This is actually a big animated thing in the, yep, I'll ignore that. Okay, there's the static version of it – no real big surprises there. The signer, of course the bottleneck in this whole system performance-wise is the rate at which the HSMs can perform the signing operations. So the first big performance gain we have there is by doing NSEC3. As we go forward we will probably wind up parallelizing operations by adding additional HSMs, but basically the signer has to cycle through different TLDs that are sending us IXFRs. So the IXFRs will accumulate sort of 1 queue per ccTLD. The signer will come around, sign that queue and spit it out, and right now the larger ones that we have, the sort of 350, 400 delegation zones are in the sort of three-, four-, five-minute range at worst. We actually have some graphing online that shows them mostly clustering around one minute.

Yeah, obvious stuff. Okay, the live demo, that was for Thursday. So okay, thank you very much.

Steve Crocker: Good, thank you.

Bill Woodcock: Sorry – Monday.

Steve Crocker: We're good. We started a little bit late, we were pressed for time and we overcame it, and now we've got the embarrassment of having a rather nice amount of time for questions or for going out. It's boring to say- No, we don't have that much time. So let me open the floor, both for this session, the signing services in particular, and I don't see any reason why we can't have a broader set of questions and discussion if we want to. We have on the clock 35 minutes, and I'm perfectly happy to bring things to a close if we're at a natural point but let's see what the discussion brings. So have at it.

Male: Hi, my name is (inaudible) IDN. I have a question for EURID and Nominet because in the presentations that they gave us and the outlines, they made a mistake or I feel it is a mistake that it's the registrars that run DNS services which in my opinion is not true. Registrars only perform administrative actions. So my question actually is are you going to do something in your administrative interface so that DNS operators can enter your interface instead of the registrars?

Simon McCalla: Yeah, really good question. Two things – what we find in the UK is the bulk of our registrars are also DNS operators as well, so you have that.

Male: That's true but in the model they aren't.

Simon McCalla: Exactly. So we will also, as well as having EPP interfaces and so forth we will also provide an online portal for people who do manage their own DNS records too to be able to access those directly. So we see it as operating for both.

Peter Janssen: And maybe some extra information there. First of all, same thing for .EU – most of the registrars are also hosting the DNS which does not change affectively, indeed there is a sizable amount of DNS hosts which are separate from the registrars. And for those they still have to pass by the registrar to actually get their DNS sets into our zone, and if you remember my slide actually there is just some extra piece of information – the IP address of the hidden master, the IP addresses of the slaves for which we send a DNS key to. And if the registrar takes that into us then we don't really care if the DNS host is actually the registrar or something separate

from the registrar.

Male: My question was actually in order to accommodate this, can you for example give the registrar an extra user ID for the DNS operator to use to interface with their system?

Peter Janssen: Okay, that's another question, and yes, if you remember the last slide and the very last line where we have that delegated key management interface that we're actually going to make available to the DNS host. And it is of course the registrar who has control over does he give it out, yes or no, some sort of a security mechanism like a password or something like that. But at that moment in time it's a DNS host that can directly interface with the registry to do the key management, but it has nothing to do with the signing service as such.

Male: And so can a registrar hand out multiple of these?

Peter Janssen: Well actually on a per domain basis we envisage for the moment as having some sort of credentials, user ID, passwords, certificates – we don't really know what yet – but you can hand that out to any party including a separate DNS host or the registrant that hands it over to the DNS host or something like that. And again, if I can continue quickly Steve, we're very interested in seeing parties that actually want to work with us to have some sort of a standardization process in place there, that actually that interface that's given out to DNS hosts, registrants and so on is actually something that would be the same for any registry out there in the world. Thank you.

Steve Crocker: Let me express personal thanks for you raising that question. The distinction between a DNS operator and a registrar who happens to be offering DNS services as part of it is one that's bedeviled and caused confusion at many points along the way. In the ICANN environment there are accredited registrars and contractual rules governing the behavior of the registrar, but it only governs the administrative transactions as you had alluded to.

It does not cover the DNS operations, and so we've found when we've tried to entangle the technical aspects of the subtleties of how do you transfer DNS operations for a signed zone for example, that even the terminology has been confusing because we needed to untangle it. So I'm very, very pleased that you asked that question. As you see you got good answers from knowledgeable people here but I think it's going to require just a little more attention and repeat activity to bring this up the level where everybody has that clarity of understanding. So thank you. Simon.

Simon McCalla: Just to echo Peter's comments, and it's a plea really to anybody here on this panel as well as to anybody, the rest of you who are in any way thinking of launching DNSSEC signing services. I think standardization is really important as much as we can. We want this to be a place where people can use multiple signing services if they wish or they can skip between different people's signing services. We see having a flexible market for this being really important to its adoption, so if it's a difficult job to move from one person's signing service to another because you've got to recraft all of your systems then that's not a great story. So we're very keen at looking at trying to standardize as much as possible the way in which we interact with these services, whether it's through EPP or whether it's through web interfaces, or whatever. So yeah, very much.

Steve Crocker: Would it be helpful to you if we started to impose some standardization rules from ICANN? No, never mind. (laughter)

A joke for those who understood it, never mind if you didn't. Other comments, other thoughts? Any other topics that you all want to bring up for discussion? Anybody on the panel? Oh boy, people are eager to get out of here.

Jim Galvin: Question? Okay. So if there are no questions maybe it's because you all already have signed zones or you know exactly what you're going to do to sign your zone? Maybe we could explore that space a little bit, Steve, with the participants here?

Steve Crocker: Oh, interesting idea. Alright, audience participation time. How many of you have signed your zones? (gasp)

Jim Galvin: Don't forget to raise your hand.

Steve Crocker: Yeah, yeah. Well of course we've signed our zones. How many of you plan to sign your zones? Good.

Jim Galvin: I wanted to raise my hand again because I wanted to sign future ones.

Steve Crocker: How many of you do not plan to sign your zone? (laughter) How many of you have no idea what a signed zone is? (laughter)

Jim Galvin: Yeah, right.

Steve Crocker: Oliver back there, whose Co-Chair of the DNS Extensions Working Group and in is public service job works for me, at least I pay him during his day job, and I can tell you for sure he has no idea what a signed zone… No, oh – and Peter, too. Good. Wes?

Wes Hardaker: Wes Hardaker, (inaudible). I sign my own zones because there wasn't any signing services when I got started. Mine runs out of a make file; I sign a number of personal and a number of work-related zones all managed under one belt, and it works just fine. That being said, I have a number of zones that would be really bad to go dark, and if I was to consider not doing it myself anymore – which I'm obviously capable of doing it technically because I've been doing it for awhile – it would be because from a management perspective, and this has been alluded to on a number of slide sets.

I would prefer somebody that has a very large background in managing data, making sure that it's good, making sure that the DS record never gets published to org when the child is you know, no longer has a key for that – not that I did that last week. That's the real service that I see in registries and registrars and DNS signing services out there, is the ability to make sure that my zone doesn't go dark in the same way that people do it to get you know, propagated DNS all over the world and to make sure that their DNS servers never go dark.

It’s going to be the same thing for DNSSEC. Yes, I can do it myself, I wrote some of the tools that help do it. That doesn’t mean that my little tiny box is quite the right production place for large zones.

Steve Crocker: I was going to, thank you very much for that. I was going to suggest, I was going to ask for ideas to include in the next DNSSEC Workshop in Singapore or the one after that. But and so I do hereby asking for that, but in addition let me ask more specifically – would it be interesting to try to have a panel on things not to do or mistakes we have made, sort of an accumulation of hard lessons of the little things that could go wrong that bite you; and have those aggregated in some experience from the early operators?

So that’s one thought, and I’m interested in responses either on that or the more general question of what topics would be useful and productive to bring up in this section.

Paul Wouters: Hi, Paul Wouters, (inaudible). I had a question for the people running signing services. Are there any provisions in place to avoid coercion of revealing your private key to your local government? Are you thinking about jurisdictions of where to put these keys or can you assure me that a signer service is in a certain jurisdiction?

Peter Janssen: Me first. In our case, local government is a special kind of thing because it's the European Commission so it's something a little bit weird, I would say. The real answer is we're in an improved concept phase. We haven't been there yet. We actually want to put something out where we actually want to see it work and work with our registrars to see what they have to say about it. And at that moment in time or a little bit before that probably we'll start seeing the whole nitty gritty details, also one of the things you raised now but we're not there yet. First let's get the technical side running and then see if it is actually worthwhile continuing on this path. That's the general idea.

Steve Crocker: I think it's a very helpful question to have on the table, but I think it would be also helpful to have a degree of clarification. So it's my understanding, and I seek responses if I've got the wrong picture, that when one uses the modern hardware signing modules, that compromising the private key is essentially impossible in the common sense notion of compromise, of exposing it to air and having somebody run off and use it in another box and sign things.

So part of my question is validation – is that correct? And then the other aspect is that the coercion question then I think becomes transformed into could a government or somebody else coerce the signing service to inappropriately sign something that shouldn't have been signed and sort of get the same effect. And so the question does not go away just because there's hardware protection but it gets transformed into analyzing where the points of entry or points of attack might be.

Bill Woodcock: So that's one reason why we're doing two separate parallel signing processes, one in Zurich and one in San Jose so that we can have obvious to the public two processes under separate jurisdictions, each of which should be producing exactly the same result.

[background conversation]

Bill Woodcock: Yes, yes there's code around that.

Male: So I'll offer a similar response to Bill. We have hardware modules, we have not had to deal with a question of different jurisdictions wanting private keys but the immediate reaction would be one of "You simply can't have it, there's no way to get at it because the box would simply make it go away." So it would be interesting to see that tested at some point.

And then like Bill, we have facilities in you know, multiple places, signers are in multiple places so that's also an additional factor in dealing with this. And those boxes that we use have synchronization facilities built into them, I mean even Matt said earlier, in the HSMs that they use – I'll speak for that – about you store encrypted blobs which only the boxes can deal with, and then you can move those blobs around but you can't get access to them. So that's how you hide the stuff.

Steve Crocker: Good. Japp? Say your name first.

Japp Akkerhuis: I feel that the problem is a little bit similar as the RIRs which are doing that with IPKI and offer to house the keys and all that stuff. I mean it's more generic to blame, I mean to source this out to a third party you have to do some of this analyzing about what the parties will be doing. I mean when it's under different restrictions than yourself it's something just to be part of the discussion before you start to use this service.

Steve Crocker: Thank you. Paul?

Paul Wouters: Just a clarification again. Remember that for signing services you have the unsigned zones so you don't need to have a compromised key or a private- You don't need to have access to the private key to just change the unsigned zone and then let it sign in the regular process. So there's no need for access to any private HSM keys.

Jim Galvin: Right, I think that's comparable to the point that I was making that you could sort of misuse the service…support the service by putting in data that shouldn't be signed if you can get access to the process there.

Paul Wouters: Right, but it also means that for instance all these signing services can sign a different version of the zone they can then give to someone else to use somewhere on the leaf, so this doesn't necessarily have to be published worldwide. You can sign a private copy with different data.

Jim Galvin: Right.

Craig Spiezle: Craig Spiezle, Online Trust Alliance. This is more of a comment and an observation, and it's really intended for people outside of this room and really gets down to… When I speak to you about DNSSEC it's really trying to articulate what is the business value proposition? Why should a business care? And so we talk about a lot of technical issues here but I think that's an opportunity for us to really articulate and as a group kind of resonate and repeat and get that and communicate. So I'd like to know your thoughts and if any actions to really help out there and really market DNSSEC effectively.

Steve Crocker: The question is exactly to whom DNSSEC is being marketed. I think that's another part of that question.

Jim Galvin: I just want to add a comment, too. I think we heard part of the answer this morning when we were talking about Mozilla and Fedora, and so putting the DNSSEC… I think it needs to be available, it needs to be incorporated. The validation process needs to be available to you to help promote the need to sign it. So you have this double-edged sword of things that need to happen in parallel, so I think that's also part of the answer.

Russ Mundy: Yeah, one of the challenges that DNSSEC has had and the technical community in particular has struggled with for quite awhile is how do you describe a value in particular of the correctness of a name? And some organizations who have had bad things happen with their name, whether it was through a hijack or something in the whole registration process where they lost control of their name – organizations like that sometimes completely fall off the internet functionally and only then do they seem to really recognize the value of the name.

And so as Steve mentioned a little bit ago it really oftentimes does depend upon who it is being marketed towards, because a name and a presence on the internet is important in varying degrees but it's also important in different ways to various organizations. And so as you, if you think about DNS being really an infrastructure piece that's historically kind of been buried down, you actually see the name but you don't as a user recognize what goes on and how important it is.

Now as people start to attack names and DNSSEC is a counter to that, those attacks are what has to be envisioned, described in a way that's of value to those being used. And there's not really a simple answer. One of the things I think we've tried to do is identify some communities that are of particular interest and now the work that's been referred to a couple of times today, the [Dane] work in the IETF, it looks like that may also produce some easily identifiable advantages of making use of DNSSEC for multiple and different business opportunities.

Craig Spiezle: Yeah, I might- I know PayPal's been doing some great work in this area, and finding those North Stars so to speak, those early adaptors… And we had this, we worked on EV/SSL certificates, similar cases – the chicken and the egg, Steve Crocker talked earlier about [Deacon], same thing as well – if no one's checking what's the value of signing? So I think there's just other areas and I'm certainly happy to work with anyone in that area of really trying to promote that as we go out there. As the .com is signed later this month we have an opportunity to do so.

Steve Crocker: So let me ask you one question – if you could be so kind as to give some thought to is there value of having a panel at the next Singapore meeting that would try to address some of these issues? And do you have any suggestions in that space – not necessarily at this moment but if you could think about it and pass it to us that would be great.

Craig Spiezle: I actually suggest we have an opportunity to do it before Singapore as we get out there and get a message, and maybe some work like some of us did with APWG and other groups, try to synthesize what that message is. Cause again there's a lot of confusion as I talk to security professionals outside – what is it? When is it going to be? Why should I care and where does it stack, rank in the priorities from a security perspective? So thank you.

Steve Crocker: Thank you. I think we have come to the appropriate time to give ourselves a collective round of applause. I thank everybody for participating in this extended session here and we're in, we're rolling right along with DNSSEC becoming a major part of the landscape. Thank you all.

[End of Transcript]