TORONTO – Tech Day 2 in Cooperation with OARC
Monday, October 15, 2012 – 11:00 to 17:00
ICANN - Toronto, Canada

EBERHARD LISSE: We're not doing the IPSec tutorial, but Paul Vixie will speak about Rate Limiting in the DNS. Otherwise the topics are as listed on the Agenda. We have this time looked for the DSSA and what was the other group again? I hate these Acronyms. SSAC — I like these acronyms. Since the ccNSO council has on the last meeting given us the mandate and instructed us to sort of see whether we cannot widen the target audience a little bit, we have started to communicate with other groups that are in other constituencies.

There is also a Cross Constituency Working Group to sort of lessen the digital divide; we also are talking to them. Our presentations or our topics are not just totally valid for ccTLDs, but also for other TLDs, in particular smaller TLDs and gTLDs. Most of the new ones will be smaller, at least in the beginning. The problems that we have solved or that we haven't solved are the same that we are facing, so I think this is a very good idea.

The first topic will be Chris Davis speaking about Abuse Mitigation the Secure Domain Foundation. He has been speaking about this in San Jose, if you all recall the ones who were there, so he will start on this. The second will be Architelos they will do this from a commercial aspect. They have a commercial product we'll demonstrate, and then we will see that we can get Garth Miller who's currently in New Zealand

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.

to do a remote presentation about how to integrate this with CoCCA Tools.

CHRIS DAVIS: Can everybody hear me okay? Okay, great.

[background conversation]

CHRIS DAVIS: For those of you who don't know me my name is Chris Davis. I'm a security guy that dabbles in DNS, not a DNS guy that dabbles in security, so I apologize if I get some terms wrong with bailiwick and things that I don't really understand.

To giving you a really quick background on myself, I started doing information security work around 1995. And I started in the pen-testing vulnerability assessment arena. Malware, viruses, that sort of thing wasn't really that interesting to me. In 2006 I was working for Dell in Austin, Texas running their Global Information Security Assurance team and I really hated that job, it was just not at all what I wanted to be doing.

And I met this really nice gentleman named David Dagon, who had a startup in Atlanta called Damballa. And he said "Hey, why don't you come over here and do this work with me?" And I was like, "That's malware work, eh," so I took the job and fell in love with that facet of security and I've been doing it ever since.

So the Security Domain Foundation is kind of the evolution that start and the friends that I've made on that journey. We're going to talk really about what the Security Domain Foundation is. The first thing you need to know, this is a 100% public benefit, nonprofit gang of folks. In no way do we sell data ever; we don't charge for our services in any way.

The Security Domain Foundation's made up of about 27 different volunteers across many different disciplines, some of the top security researchers in the world. You would know their names, but I didn't ask them if I could name-drop them so if you really want to know, you can find me afterwards. We have researchers that use our API and our back-end data. And those researchers come from places like Facebook and Google, and Microsoft, and Trend Micro, and Kaspersky, and Semantic, and MacAfee, and on, and on, and on.

When I was in Costa Rica we pitched the idea of creating this foundation and our original idea there was that we had this malware analysis infrastructure, which we thought was pretty cool. And we were analyzing all these volumes of malware and pulling out what that malware was talking to on the internet and how it was communicating. And I thought, you know, what we could do is we could generate like a daily list of bad domains on a per registrar or per registry basis and just provide it to those people and everything would be wonderful.

And so that was the original idea. It didn't happen. And the reason it didn't happen is that a lot of registrars don't really want to go back and deal with things that are in their space; they're kind of apathetic to it. It costs them money and time to go suspend domains after… If they don't

have an abuse report in front of them, they're probably not necessarily going to jump right on it.

The other thing we noticed is that it's actually really hard to generate lists of malicious domains without having any false positives in it. And just the way that malware functions is it looks up five or six different domains and only two of them are actually malicious, plus maybe one of those two is a compromise site, versus registered for that purpose. So we very quickly sort of got mired in trying to come up with ways around that.

So we changed our path and the path that we've changed to is a JSON based API where people can query many different kinds of queries and get back lots of interesting data. We don't really have much of a public face. We have website at thesecuredomain.com and .ORG, which bounces to .COM. And it's just one-pager describing our mission; the public front is coming soon.

So, what are the issues? And I think we're all pretty clear on this, but there's a lot of recidivism in abuse. Bad guys get their domain taken away at one registrar, they get their account suspended, and then they go to another registrar and they create another account and register other domains. Sometimes in fact the way that the handling of abuse occurs is that the bad guy will just simply move his domain to another registrar. I've seen that happen as well.

So the bad guys don't give up; they just become somebody else's problem. And what's interesting is that bad guys, their WHOIS data is all B.S. — we know that, or it's mostly all B.S. But they use the same B.S. over and over again and I found that pretty funny. And I think the

reason that they do that is because they don't have any need to make up new B.S.; there's no sharing of this information between different registrars or between registries.

So I'll have an email that I set up with Gmail or Hotmail or something and I've got some B.S. information I put in and then I get suspended and to a different registrar and I use the same email address. So it becomes actually fairly easy to track these guys because they do the same thing over and over again. Of course privacy protect makes our life difficult, but I understand why it's there.

So there's really no incentive either for registrars or registries to share this abuse data between them outside of altruism or the public good. But if we can facilitate that data sharing then you can not only prevent others from inheriting your headache, but others can help you from inheriting theirs. So that's basically what the goal of the Foundation is.

What do we have today for data at the Security Domain Foundation? Well, we track on a permitted basis over 260,000 unique bad individuals. We collect lots and lots of WHOIS data, however we only started collecting the data going back to mid-February. So we have almost 26 million WHOIS records in our dataset and it updates every day.

One of the reasons we didn't go after a lot of historical data in WHOIS is 1) it's really expensive to get domain tools. And 2) the bad guys don't tend to use the same domain for longer than a year. They don't register it and then keep renewing it, unless of course they're completely undetected and they just keep their botnet rolling. But we find that

statistically most of them keep the domain register for one year. If the domain stays live for a year, they'll move it to something else.

We analyze lots and lots of malware through our partnership with Emerging Threats, which is one of the founders of the Foundation as well. So we currently have over 5 million malware samples that we've analyzed, that we've put into the database with this malware md5; looked up this domain; this domain is categorized this way and here's the date, etc. So the growth rate there, which is 100,000 a day, that's on average the number of malware samples that we process every day.

We've integrated most of the public lists into this. I'll talk a little bit later about the collective intelligence framework from the REN-ISAC guys, Wes gave, but most of the major domains. And then I also have a small for profit startup that I'm not really talking too much about today. But it has some really interesting data that it donates to the foundation as well. So I'm going to talk a little bit about that unique data that we donate to the Foundation. Okay, so I can't tell you what the data sources are, but if you're really super interested and you want to sign an NDA we can talk about it.

So in just 90 days of operation, which is how long my little startup's been around, we've got 2 + million logins from known bad guys, where they've logged into an account, updated something, be it hosting or domain A records or whatever. We've identified 96 thousand new accounts that were created by known malicious bad actors. We've collected over 163 thousand browser fingerprints. Now, if you've ever been to EFF's Panopticlick webpage, it describes the browser fingerprinting technique that we use.

But what's interesting is if you just give a browser a little bit of JavaScript to run it can query for things like: What's your screen resolution? What language packs do you have installed? What browser plug-ins do you have installed? What font packs do you have installed? And all that information is replied back, and when you take that information and calculate a hash from it, you end up with a unique fingerprint that's about 90 + percent unique. There are some collisions, but not much.

And we have a lot more than just that. So my example here is where I say okay, this data allows us to say bad guy XYZ accessed his account using this email address from this source IP 12 minutes ago. He updated an A record to point to here. He connected to a VPN service; he connected over here or updated his hosting account — that kind of stuff.

Our unclassified data pool. What I mean by that is that there's a huge pool of information and we only track the guys that we know to be bad actors. Where they've already breached the terms of service with that given provider or — again, I can't talk about the source of the data. So when I say unclassified pool I mean this is the people we don't know about. But we know that there are a lot more bad guys in there than the 260 odd thousand that we're tracking.

Okay so use cases, how could you use the Security Domain Foundation's API? Well, registrars can query — this is all pretty self-explanatory. I don't know exactly how every person would want to use it, so we take the best guess that we can. We try to provide that service and if there are use cases that we're not thinking of, let us know. And you know

what? Our turnaround is like 24 hours to give you a new way to access the data. I don't really need to say it on this slide. You guys can figure that part out.

Okay, so what's the cost? As I said, this is free and I mean it. You can't donate to get access. There's absolutely no way you can pay for this data. We only provide access to the API, to registrars, registries, hosting providers, DNS providers — basically infrastructure people, and some security researchers, and of course the volunteers at the Foundation. But this is not a "public" API. Because of the sensitivity of the information we keep it restricted to people within the infrastructure.

So how does the API work? It's pretty simple. You can query by an email address. You can query by IP, be it the source IP that you have or an A record that somebody's pointed something to. You can query by domain. Now, what's interesting, when you have a registrar that somebody for some reason is suspicious to a registrar and he wants to register crazydomain. — name your TLD.

Well, that domain is not yet registered, so it doesn't make sense to query the API by the domain. However, we have wildcarding enabled as of this morning, which means that you can take the crazy part before the TLD and just query it. Imagine it as a keyword search. And we would come back with every known malicious domain that matches that keyword, that we've ever seen. Obviously we don't have every known bad domain in the world, but we've got a pretty good start.

If you're really into security or you've got a great security team and they have a particular malware md5 that they want to query by, you'd go ahead and do that. If you happen to be using the Panopticlick style

JavaScript to do browser fingerprinting instead of using cookies to track your users, you can query by browser fingerprint. You can query by a user name or an alias, and we'll come back and say "Hey that user name was also used by this bad guy over here to do this thing." And the name server part, it's just an indexing issue we're having right now, but that'll be up very soon. I think that covers most of the ways that you can query.

Now the integration currently, CoCCA — they're going to speak a little bit. They've integrated this as of a few days ago, so it's working across I think 12 or more ccTLDs. Maltego, which is an open source intelligence tool; it's used by a lot of different security researchers. Trend Micro was nice enough to write the Maltego transforms for us, so those are available. Case File is also a product by the same guys that do Maltego. These are open source and free if you want or you can pay for a license.

Palantir is a commercial product that we're working on integration with for different research purposes. And the Collective Intelligence Framework is a really cool — I don't even know how to explain it. It's a really cool framework, written mostly by Wes Young at REN-ISAC, and if you haven't checked it out, you should. So I'm just going to just jump from this to show you kind of what the API looks like on the web browser, and then I'll take questions, and I'll be done.

So this is the Maltego tool, if you can see this all right. And this is me taking an email address and then running a transform, which is the Security Domain Foundation email lookup transform. And the response it comes back with is: okay, this email address is related to malware registering camatic control domain. This email has been seen active in

black market forums, selling credit card numbers, trading bots, things like that.

And here are the IP addresses this guy has logged in from or pointed domains to, and then you could using the Maltego tool select the IPs and okay, show me all the domains. You could use something like ISC's passive DNS database at this point, which we actually have written a transform for that we need to give to Robert and Vixie in case people want to use it that way. Okay, so that's the Maltego instance of that.

Here's what it looks like on a web browser. This the Collective Intelligence Framework where I just ran an IP address and it came back: okay this botnet infrastructure according to — okay I can't see this part of the screen — according to AlienVault, this was also searched for by somebody else in the — this is via CIF.

This is what it looks like when you use it just as is and this is my API key, which I'm going to have to revoke after the presentation. You see this is me searching by email. Well, we've got a login failure here. We know this is our bad guy for these two reasons. He's updated his account on these dates from these IPs, pointing A records to the Res or Resolve IP. Logins, updates, you kind of get the idea here; I don't have to keep going.

This is a search by domain where we want to know if this domain is bad. Well, the malware md5 count comes back at 74,720 pieces of malware have looked up that domain — it's probably bad. Actually it's funny an analyst did categorize it as compromised at one point, but CnC spam, compromised, spam CnC.

IP, this is by source IP. This email addressed logged into an account using that source IP and if you scroll down we'll see that the email address has changed. This account updated a DNS record or their account in some other way from this IP — well, it's Res IP, so they updated a DNS record. So we see the changes there; you can kind of tie people together with that. And lastly, Resolve IP, this is somebody pointed something to this from this email address, etc.

Okay? And that's it. So, questions? Yes, sir? I recognize you.

ROY ARENDS: Thank you. My name is Roy Arends. I work for Nominet. I know a little bit about this. We spoke about this in Costa Rica. And I really, really like this tool. Just out of curiosity, you're not accepting any donations?

CHRIS DAVIS: No.

ROY ARENDS: You don't have to pay for the data.

CHRIS DAVIS: That's right.

ROY ARENDS: Who pays for all of this?

CHRIS DAVIS: We do. The volunteers do.

ROY ARENDS: Great. Thank you.

EBERHARD LISSE: Any more questions? We have a little discussion afterwards when all three presentations have made. And I think that will be probably more interesting than doing it after each. The next one will be Architelos.

[background conversation]

EBERHARD LISSE: Before we start, the live browsing will not be accessible for the time being on the remote participation. But there is only three remote participants at the moment and one of them works for my city TLD, so I have just instructed him to behave himself.

GREG AARON: My name is Greg Aaron. I'm here with my colleague Michael Young. And thank you Doctor, for having us today. We're going to talk about how abuse can be detected and then mitigated in TLD spaces. First I want to define what we're actually talking about here and put some bounds around it.

Every service provider, whether it's your credit card company, your phone company, your registrar, and a lot of registries have terms of

service — a contract with its users and this is basically what regulates behavior on the internet. Now it is true that law enforcement for example, may occasionally get involved in certain cases, but that's a very tiny number.

Basically services regulate what's going on, on their networks or on their services and that's how things are done. You know, if PayPal finds someone laundering money, PayPal will shut down that account. If you do not pay your credit card bill your bank will cancel your card. And of course registrars and a lot of registries have terms of service as well and those define for their services what they find to be unacceptable. And so of course those vary a great deal amongst providers.

Today I want to talk about things though that I think almost everyone could agree are problems, and let's leave it at that for now, but things that are designed to exploit internet users — basically behaviors that can be malicious or would be considered criminal pretty much everywhere. So activities like spreading and running malware, phishing scams designed to perpetrate identity theft or theft of money from internet users, you know, running botnets and those kinds of things.

A lot of service providers will also define other kinds of things, which might also be problems for them and they don't like to see those. Things like brand infringement, cybersquatting, hate speech, those kinds of things. We're going to put those aside though, for today and concentrate more on the first set of things I mentioned.

So first, what's the landscape? As of right now there are over 240 million domains in the world's TLD registries. And there's been quite a lot of growth over the last year actually. Now a lot of those domains are

older domains; they're older than a year. .COM and .NET, for example renew at 73%. A lot of TLDs are a little lower or a little higher, like .EU for example is at 84%, but what it basically means is there's a lot of churn.

And a lot of the problems that we really need to concentrate on are with domains that are recently registered; they are less than a year old. These are the ones that bad people register and then they throw away after they're done with them. So that pool is about 73 million domains that are less than a year old, and we know that because we can look at zone files and so forth, and those numbers will match up, so about 73 million are in that pool.

Now within these groups of domains, how can you find out what's happening and what's going wrong? There are a lot of different methods. One of course is to go out and spider; that's how Google will display search results of domains that might contain malware. They go out and see what the sites throw back at their spiders. The antivirus companies have software installed on desktops and laptops and then when users encounter problems those are reported back to the security company. You can do traffic analysis; you can look at DNS query data for example, and see what's trending and from where.

And then one of the most interesting ways to do it is to look at domains that are being advertised in email, because email is the way that bad people advertise bad things. That s how phishing is advertised, that's how all the frauds are advertised, that's how links to malware are advertised. And then of course those people sometimes get roped into botnets and then the botnets send out more email.

So let's talk about that for a minute, because it's an interesting indicator of what's going on. So we're going to talk about domains advertised in the bodies of emails. I am not going to use the word 'spam'; I'm going to put that aside. Now, most people would define spam as bulk unsolicited email, but the problem is some places that's okay to send — it varies on jurisdiction. Instead what I want to concentrate on is why is mail being sent and how is it being sent. What is the purpose and what is the intent of the sender?

So if we look at those things, how and why and who, the unfortunate result — and there's a consensus about this — is that 75% to 90% of all emails sent in the world are sent for abusive purposes or in an abusive fashion. And it depends on the exact methods used and who is studying it, but all the studies fall into this range.

And that's work by security companies like Signin Tech, and then also organizations like MAAAWG, which is the Messaging Anti Abuse Working Group, which studies these issues and is the industry association dedicated to solving these issues, using very large datasets. And this has been historically the range.

Now how is most of that sent? Well, a lot of it is unfortunately sent from botnets. Most of these big botnets are rented out, sometimes to other users or sometimes to the owner of the botnet. And they're advertising and sending out mails from these zombie computers. Two of the most prolific are called Festi and Cutwail. They're literally sending billions of mails each day. Now you don't see those emails because there is whole group of companies and people protecting your mailbox,

but that is hitting your network and some of it gets through and some of it doesn't.

Recently the Grum botnet was shut down in August and the researchers were able to take apart some of the files associated with it. And they found out that there were about 2.3 billion email addresses that this botmaster had access to and was using, so probably every single one of us in this room was on that mailing list. We got it — you may not have seen the mail, but you probably got some from that botnet.

A lot of this mail is also being sent through Snowshoe spam. And it's called Snowshoe because when you wear a snowshoe you're spreading out your footprint. And what these people do is they obtain IP ranges and they usually do it by lying about who they are, sometimes they'll even highjack a block from its current owner, and then they'll send out spam throughout that block. And basically the idea is to be able to send out that spam as long as possible without being detected, and shut down, and put on a black list, so it's a very dishonest way of doing things.

And a lot of the mails if you click link you'll be sent through a redirect. The domain name that's being advertised in the mail will then send you somewhere else, to another destination. And the idea is to again, avoid getting blacklisted for as long as possible. This is one of the reasons why spammers or whatever you want to call them buy a lot of domain names, because they will always need to have fresh domain names that have not been blacklisted.

So here's kind of a typical distribution of what some of these spam campaigns are. Cutwail varies, it depends on who's using it or for what

each day. About half, sometimes as much as three quarters is basically dedicated to spreading malware or some other kind of identity theft.

On the right for example is a sample of a mail that was sent spoofing the Internal Revenue Service. This is the tax authority of the United States and it says "You are due a refund," and if you click on that link, one of two things will happen. You will get a drive-by malware or you will be taken to a site where you're asked to put in your personal information, such as your bank account number where this money can be deposited.

A lot of the mail that's advertised is what's called rogue Pharma; these are basically all that Viagra and Cialis spam that's out there. And that's run by organizations often in Eastern Europe; they're basically organized crime — it's pretty bad stuff.

So how many domains are involved? One way we can look at this is to see what's being tracked in these blacklists. One of the most important is called SURBL and this used to protect literally billions or mailboxes across the world. SURBL has about 800,000 domain names on it, on any given day and they're adding about 6,500 or so. So over the course of a year the SURBL tells me they listed 2.4 million different domain names on their lists in that 12 months. And most of these are recently registered. They're used for some purpose and then they're thrown away and new domains start to appear.

One of the other big providers is Spamhaus. And they also maintain a list of domains. There are currently about 330,000 on their list, about 4,000 a day coming onto the list and old ones being retired. Spamhuas listed almost 2 million unique domain names on their Domain Block List over the last year. They also maintain an IP list of IP addresses. They

have 6.7 [million] currently listed and almost 9 million over the course of a year.

So what does this mean, put it all together? Just looking at SURBL and Spamhaus together — they do have an overlap — but what it means is they alone listed 3.5 million domains that they recommended people not go to at all. So if you look at the 73 million that are recently registered in the last year, what that basically says is 5% of new domain name registrations end up getting listed on these lists. They're being used for bad purposes and they're usually being sent in very inappropriate fashion.

Of course 5% is high for TLDs and some TLDs have more than 5% of their new domains involved. But 5% is unfortunately probably the floor. Those are the domains we know about just with looking at those two lists. But if we start to look at other lists and of course there are always things we're not catching, we know that the number is probably above 5% worldwide. Add in extra sources of data — the block list users are certainly not catching everything — some slip through.

And then we know we're missing things for other reasons. At the Anti Phishing Working Group we started sharing data with CNNIC, which is the Chinese registry and they operate anti phishing association in China. We found out that most of the world was missing most of the phishing that was going on in China because outside of China people were not parsing and properly reading the spam mails, because they were in Chinese. But the Chinese were catching this stuff and we were seriously undercounting a lot of phishing as it turned out.

Now how many are getting suspended? This is a real unknown, and I'll stick my neck out and say I'm betting probably at least a million domain names are being suspended per year by registries, by registrars. And of course sites are getting shut down by hosting providers when they have problems reported to them. We don't have a firm grasp of this number because those companies don't tend to talk about their operations for various reasons. But we can see some of this activity taking place in zone files, because for instance registrars will point domains to certain suspension name servers when there's abuse.

So it's a reminder that there is a lot of activity going on and suspensions are actually a very routine part of what goes on in the domain name space every day. This is a good thing, but it also is an indicator that a lot problems are not being taken care of. Unfortunately abusers can consume a lot of domains. Again, once they get blacklisted they'll need to move on to a new set of domains, and they'll obtain them in various ways. But they tend to be evasive; they want to kind of keep below the radar.

One way they'll buy the domains is by using gift cards. Now, these are issued by companies like Visa and MasterCard and they have it's basically a credit card number on the card, right? But it's not associated with someone's name. You can buy it for cash and then you can use it online to make a purchase. That's a very popular way of doing it, because it's never associated with an individual. And it's also a good way to use laundered money that's gotten through some illegitimate fashion.

A lot of these domains will use falsified WHOIS data. Sometimes the data's made up, but criminals are getting pretty good about using legitimate looking information. They have access to large numbers of legitimate names and addresses. So if you look it up and try to validate it — it's a real address, it's actually a real person, but they're not the one who bought the domain name.

There are certain places that are known to be good places for buying domain names because the operators tend to look the other way. There are some resellers that have issues and there are at least three ICANN accredited registrars who are owned by spamming organizations, unfortunately. So there's going to be a supply. The question is where are they going to get these domains and in what TLDs? And these people do tend to move around. They'll go where they can get things that they need.

Michael's going to actually show you a little something.

MICHAEL YOUNG: Thanks Greg. Just take a little second, because I know a lot of people in this room may not know who I am, although Greg's well known through his work in the APWG, by most of you and other presentations he's done in front of you. I've worked in the industry now since about 2001 and I ran the technology and the operations for Afilias for a lot of years before joining Architelos as a CTO just over a year and a half ago.

We struggled in Afilias. Like everyone, we struggled with managing these types of issues effectively and cost effectively and judiciously with good tracking and good proofs. It was burgeoning space when we first

started to wade into this. Greg was working with me at Afilias at the time and there was a lot of learning that happened back in 2005 – 2006 when we started to really wade into this problem.

One of the things that we've been able to realize working together at Architelos is a tool that helps people work through these issues or work through these problems in consolidated fashion with some automation. One of the biggest problems is just the overwhelming data that you're facing. If you want to look at multiple sources and you want to take action on some of these issues as an operator, or a registrar, or another industry stakeholder, it becomes very, very challenging because of the sheer volumes. So we developed a tool called NameSentry to assist with that, so that's what I'm about to show you.

Now what you're going to see is an administrative view when I first log in that is absolutely the entire TLD space in the world across a number of data sources. You can see them down here. A couple of them Greg's already talked about, SURBL and Spamhaus. You have Internet Identity down here as well and a bunch of other ones there that probably look familiar to most of you who have been involved in this space. The tool consolidates and brings in those data pieces and brings them into a common format, so that we can digest them, compare what's being reported on the sources and provide a unified view to the product.

So what I'm going to do is I'm going to dive down into a couple of the TLD views using gTLDs, because I think nobody wants to go in-depth into their ccTLD in front of everyone right now.

MALE: You can do mine.

MICHAEL YOUNG: Alright. But we'll start with .COM anyways — our favorite space to take a look at for abuse for everybody, because of the size of it and the sheer volume. So you can see 421,000 active abuses. This means in this case these are things that are currently on these abuse lists that we have feeding into here and the sources. These are current, so that is as you can see a very large number with 8,000 of them in the last 24 hours.

This is your landing page as a TLD view, but you immediately go into some quick links to dive down into details. So you can see you can jump down by source, if you're particularly orientated on the source, because these sources have different characteristics and different policies. So your impression of them may be different and the actions you want to take on what's reported might be different, given what source you know that the report's coming from.

So I'm going to dive down into abuse type — that's another way that we sort things out — and look and see if we can find any interesting phishes. I was looking through this just before the presentation started. This data is coming in constantly. I'm going to see if I can get the whole screen on here. I'm seeing if I can find one that we saw just a few minutes ago. There it is. BDFacebook; Greg and I were looking at this just before we came up because we thought it was a really good example.

So when we dive down into here you see a view that will let you know these are the various data sources. And if there are multiple data

sources reporting, it'll contain a snapshot of that data. Some of these data sources allow you to drill down further at this point. And of course when you look at the domain or URL being reported, we do a WHOIS lookup as well at that point in time. We're careful with our WHOIS lookups because we don't want to be accused of mining anybody.

So let me show you; we took a look at this and look what happens when I try and click on this because we don't want you to go through a threat, unless you really want to cause yourself some trouble. I'm real clumsy on Greg's laptop because I'm mostly a Mac guy — funny how that is huh? We'll take a look here and go down and… This site when we looked at it a few minutes ago was really impressive, actually.

Does that not look real? So this has been up just over seven or eight hours ago, and it's still active, and still causing a lot of trouble. I don't know if you can see it on the screen, the detail might be too fine, but these links to the iPhone app and the Android app are actually genuine links to those apps. So they've copied the site to that degree.

Now, once I've decided I want to do something with this, once I've discovered this information and consolidated it, and I'm starting to track it, the next thing I want to do is try and, based on my policies, create some automation if I can. Because the more automation I have, the lower my cost advantage in these abuse issues are, no matter what business I'm in, be it a registrar or a registry or another stakeholder. So this tool has an open framework to create workflow rules and actions.

And in this particular view of it I can show you a couple that are already here in this one, although you can create almost anything that you needed to. There's a quick link under Actions to generate an example

email that might go out to your registrar, if you're a registry operator. And I think when you saw on the Details page, if I go back here, you'll notice that there is an empty column here for registrar.

And that's because when you sign up with the tool and you're a registry operator you have an API to upload consistent and privatized data about the registrar's associations with your domains. So you can use that to make the data information and tracking retro… In other words, you can start doing things like tracking abuse behaviors by your registrars and seeing who are the really great registrars that are keeping the abuse down, and who are maybe not doing the best job.

So let me show you some of the other automation rules. In the panel here we have the ability to create a rule. And I've gotten a sample rule already created, because I wanted to select a couple of non ccTLDs, just to be gentle with everybody. So I picked .INFO and .BIZ here and I decided I want to track phishing, and I created a Priority Queue. That means basically I created a view on .INFO and .BIZ and I want to be focused just on phishing. So you'll see here now, I have a mixed view of .BIZ and .INFO names and phishing sites.

If I go back and I try and edit this rule — and you're going to see it's going to look a little messy because I'm an administrative level account; it's a little cleaner in the actual users. But I want to drop down here and show you when you create these rules, you can select the data piece that you're interested in because like I said, some of the data pieces have different characteristics.

For example, Internet Identity does a hand verification of all their phishing listings. For that particular data source, people often have a

degree of confidence that they don't need to do any subsequent verification on that report. And so they may decide to, if it's an i.i.d. and a phishing report they may decide to go down here and immediately send automatic action to send an email notice to the domain's sponsoring registrar and open a ticket in their ticketing system so they can start tracking the incident.

Any questions on that? Does it make sense to everybody? Okay. Now some of the questions I've had before, when people had seen this tool in its earliest versions, the first thing they said to us was, "What about IP addresses?" Because it's domain centric, it starts with the reports on the domains. So I'm actually going to go against everything that they tell you not to do about product development and I'm going to log into a beta system and show you a little something that we've started to work on — just an example.

So as all these many, many abuse reports come in we are actually mining DNS information on them. And so underneath all this is a big bad database of related IP MX records, name servers. Generally useful information that's associated with the vectors of these abuses being reported. And so in this beta version I started to take the underlying data and started to create some usefulness around it.

So you'll see this view, again, it's not quite as complete as the last one because it is a beta system, but you'll see a note for IP monitoring here. And that's one of the first things we did with this, is to allow myself to — especially if I'm a hosting provider, this is a nice feature — I want to know if I hand out some IPs to my constituents, my customers that

they're not going to get my whole block listed on SURBL, or Spamhaus, or someone else. And if they do, I want to know about it right away.

So what you can see here is I created some rules to track different IP ranges, so you're notified. The action here is to email and that's email and operations desk; it can also SMS an operations desk. But what it does, it lets you know right away that your customer basically has taken that IP range that you've given them and they're being naughty, so you can take some action.

GREG AARON: Yeah, so this is one of the things you can do. If you're operating your own infrastructure, you can see what's going on inside of it. You can monitor what's going on at your hosting provider, seeing if adjacent blocks have problems on them. You can also start to see if there are botted IPs where a machine's infected, and those kinds of things will pop up notifications to you.

MICHAEL YOUNG: So we're running out of time, so I'm going to log out of here and we'll pop back to our presentation just to close out.

GREG AARON: Okay. So to kind of go back to the high level, let's talk just briefly about what's happening in the new TLD program, which is going to change the landscape a little bit. It's certainly going to create certain spaces that have more troubles and there'll be spaces that I'm sure will remain very clean and relatively unaffected by abuse.

In the application process ICANN gave some incentives for applicants to be more proactive and also have terms of service that address some of these kinds of issues. And it turns out that most of the applicants are really interested in that kind of thing. They're proposing takedown programs where the registry operator will report things to its registrars, but the registry operator will also be willing and able to suspend domain names when it feels it's appropriate.

And then some applicants are proposing proactive monitoring and mitigation programs. Some of those are general interest TLDs and then some of them have very specific focuses where a higher level of security is probably desirable. But across the board everybody's kind of jumping into being able to monitor and then take care of problems in their spaces. So this is I think going to change the space a little bit.

So as registry operators, what would it mean for you? As I said the landscape's changing. The competitive landscape's always changing. The new TLDs will certainly introduce some new options for users. The regulatory landscape is also changing. A lot of ccTLDs of course have obligations to their sponsoring organization and or their governments.

And in general it seems to be the case that governments are getting more and more interested in these kinds of issues — interested in crime in general, cyber security and so forth. And we're seeing in some TLDs change their policies because of government interest. And that is in general a good thing.

And of course the crime landscape is always changing. Crime is a business above all. The people who do it are in it to make as much money as they can, so they're always becoming more sophisticated in

what they're doing. And it is a task to keep up with them and what they're doing.

So you need to think about, I would suggest, the needs of your organization. All of you are in unique positions as far as governance and so forth, but thinking about the risks to your organization and to your users. The risks of doing something and also the risks of not doing anything at all are significant.

And then you have to craft policies and procedures that are right for your situation, but they also, I would suggest need to be effective. And you need to come up with, at a minimum, a situational awareness to understand when things can affect you, when things can affect your brand, when things could even affect your own infrastructure.

So we're coming to the end of our presentation. We did announce last week there will be a lot of new TLD operators using the NameSentry system, those include these three here. And what we'd like to do is open it up for any questions that you might have.

[background conversation]

EBERHARD LISSE: Any questions for now? We have of course a discussion after the next presentation, that Stephen Deerhake is going to moderate.

GREG AARON: Paul's walked up.

PAUL VIXIE: Gentlemen, thank you for your presentation. I'm Paul Vixie, ISC. I'm concerned as I watched these presentations about the cycle time that you have in your heads. There's a picture in your heads about how long this can take — the process of discovering that something is bad and taking it down. You have about a half an hour from the time the domain is created until the time that they have made as much money as they need to make, and it won't matter to them whether you kill it.

So if you're proposing something that could take, as currently takes, 72 hours to cycle through all the different approvals, you are not going to be relevant in this space. The reason that we have created RPZ is because it's the only way that we can get this stuff down in five seconds. I know you guys can't do five seconds, but please tell me the order of magnitude you were thinking of?

GREG AARON: Okay, thanks Paul. Yeah, Paul brings up a point that once starts to happen then how long does it take, because criminals know that they have a window of opportunity. Now, one of the things you need to keep in mind is there actually a lot of times is a gap between the point where domain is purchased and when it is actually used. And it is usually not 30 minutes; it sometimes weeks or sometimes even months.

The people who buy large numbers of domains will by a portfolio and actually work their way through it. So when you get a first indication that a portfolio looks bad you have an opportunity to take out the entire portfolio, which is a proactive way of doing things and it actually prevents some damage from being done. Long term it also has an effect

of trying to chase this person out of your TLD space and I can attest from personal experience that that works.

Now phishing for example or a malware run of email, the mail goes out, but then the mails get clicked on over a period of time. You want to catch those as soon as possible and this system is bringing up the alerts as soon as they're detected by various sources, which are actually capturing those emails and then bringing it out real time, so time sensitivity is important.

Some registries are going to have different policies about how they want to handle things. Some really want to report things out to the registrars and give them a shot at first; that's a courtesy sometimes they extend to their registrars. It may allow some damage to be done in the intervening time though, and that can be a problem. One of the things that we're doing is that you can prioritize alerts depending on what kind of abuse it is, so if there are certain things that you feel more important or more timely, you can handle those in a different fashion.

MICHAEL YOUNG: So I want to just add to what Greg's saying here. There are a couple of points of consideration that we always balance with this from a business and a policy model that leads into good technological tools. And one of the big balances is that you don't want to take down the wrong domain, because that opens you up to liability issues. So that was one of the reasons that we worked with multiple data feeds. We felt a tool consolidating a lot of sources was important.

We also bring in the related DNS data, which allows us to do an advanced level. You didn't see it in this version of the tool, but we have been working on it in our labs — we are putting together a very, very advanced heuristical pattern matching that allows use of what Greg said.

For example somebody might register portfolio domains and then light up 10% of them in a campaign. Well, what we're able to see — because we're doing multiple feeds and mining DNS data, not just once, but continually against the things that are reported — we can start to see patterns on common name servers, common subnets. We can start to get to the point where we're cross-alerting different registrars or even different registries that they're part of a larger portfolio of domain names that are waiting to be triggered.

So this is where some of the complexities come in and what we're trying to do is roll that up into something that's digestible and actionable by people in regards to their policies.

GREG AARON: So we'll stick around for the panel and thank you very much for your time.

EBERHARD LISSE: Thank you. So the next presentation is Garth Miller, who is remote. Can you switch him on?

MALE: I'm going to have to have a couple of minutes to get the bridge up.

EBERHARD LISSE: Okay. Thank you.

[technical difficulties]

EBERHARD LISSE: Personally I'm always a little upset when on a technical working day the technology doesn't work. But fortunately we have a backup. [Gillian Morris] can do the presentation. We must just connect her properly.

Garth is on the Adobe Connect so he can listen in. He is on the Skype so he can do this. Garth can you hear us? He doesn't seem to be on very much. Okay, I'm not going to wait longer. We are going to start this now from here, and if you get him online he can take over anytime. Okay? Especially during the panel discussion, it's probably good if he's on the panel discussion; that is even more important than making the presentation. Okay, Gillian, you're welcome.

[GILLIAN MORRIS]: Is this on? Okay. I would just explain the new feature that CoCCA has recently rolled out. So this is working in conjunction with Secure Domain. The primary goal as it says here, is to identify and take action against domains that violate the TLD's Acceptable Use Policy. And the secondary goal is to identify and share information on other domains registered by those individuals for closer inspection. Obviously people have privacy concerns. People don't generally like having their websites

scanned. They also worry about sort of innocent users being caught up in this, but CoCCA only uses publicly available data.

So the previous solution offered by CoCCA was to use a security company and do periodic scanning. This was very expensive and time consuming and ultimately not very useful. Most violations and most criminal activity takes place at the lower level, so domains don't appear in the registry and weren't scanned. You also can't get further information on the data — it's difficult to extract and very time consuming. [It] also tended to get a lot of customer support problems.

So the new system enables you to compare domains and hosts and emails in the registry against databases that contain data about these harmful users or bad actors. And so it's very simple to set up. You configure your CoCCA account to connect to an external database via an API. This does continue all the walkthroughs in the background to look for matches and as soon as one is found an administrator will get an SMS.

Crucially you can also configure an automatic response to this. So this is really, really important, because this kind of automation allows even very, very small operators with limited resources to take very proactive and time efficient measures to identify and react to the threats. Future versions, also we're hoping to have deeper responses to different kinds of variables.

So just a screen shot to very simply explain how this works. You get your credentials from the information provider and add them to CoCCA. So the information provided in this case is Secure Domains, who we saw earlier. Then you configure the desired actions, lock/suspend/exclude.

You can immediately block the domain name, so this is getting down to the response times that people were talking about earlier. This can be done almost instantaneously. Also configure your notifications for when the administrator will hear about it, either by SMS or by email. And also any further action that might need to be taken, like notifying law enforcement.

So this is the record of the data — very self explanatory. There would be a lot more history as time goes by and as records are accumulated. So just in summary CoCCA stores all unique data reports, associates it with a domain's history. It can also clear domains if the issues are proved to be immaterial. Allow drilldown for all domains registered or using the same email or hosts. Send notices to admins if a new domain is registered by an individual who's listed as a contact for a domain that has been flagged.

And just a quick summary of CoCAA, it's an association of top level domain managers who share the expenses related to development and maintenance of software. It provides the software free to everyone, but it's funded by the users via support contracts, and complies with all gTLD standards, and is used by almost 40 TLDs. 12 are hosted at the CoCCA Data Center and as of Friday this feature is up and running on the 12 that are hosted at the Cocca Data Center. And it will be on all CoCCA supported TLDs when they do their next software upgrade.

And that's it. And for more information you can contact Garth at the email address below. And I think we might have a live connect now as well.

EBERHARD LISSE: Okay. Thank you very much. I'm so thankful that you made your way in here and did it and did it so ably. We're still encountering a bit of technical issues I understand. Oh, I have here Garth, now that he can speak remotely.

[background conversation]

EBERHARD LISSE: I told him to dial back in, but anyway. Let's go and start with the panel discussion. I just wanted to say that NA runs the system on their own hardware and I usually wait ten to 30 days for a new version to be in production on the Data Center in Sydney before I upload it. That gives me a little bit of time to discuss, to get credentialed with Domain Seeker.

Because even if we're too expensive for bad guy to invent a Namibian domain name just to get it taken away, I still want to be able to say we use it and to be able to test it. And if I have somebody who is not behaving himself I can go and have a word with them. Fortunately the Namibian domain space is so small that they all know me and they all fear my wrath, which is probably worse than if I report them to the police.

Okay, Stephen Deerhake is going to moderate the panel discussion. We've got two microphones. We'll probably first have a go at the presenters and then the floor is open. We are not pressed for time; we are ten minutes ahead, so we can discuss at length.

A small housekeeping announcement — this time there will be no packed lunch. I have to say that it was extremely expensive, so we decided not to waste sponsor money on this. And therefore, I waited with the announcement shortly before, so that nobody ran away from it.

STEPHEN DEERHAKE: Thank you Doctor. I'm Stephen Deerhake. I manage .AS, American Samoa, as some of you may know. First of all I want to thank Chris and Greg and Michael for their presentations, as well as Gillian for stepping in for Garth.

My major takeaway from what we saw here is that — and this is reinforced by Paul Vixie's comment that it's a react mode. And we're still looking at like 5% of all new registrations being those by bad actors, and given the number over a month, that's a lot of registrations.

And what I would like to ask the panelists is if they in their studies of this problem and in developing these products that have similar goals — but are coming at it slightly differently — is to what extent they have thought about being proactive in looking at the vast amounts of data they have?

And doing statistical analyses on such issues as frequency of Zone File updates by the TLDs, the entry price point of the TLD. In the cc space it ranges from zero to quite expensive for example — minimum registration period is part of that as well. And I just would like to hear the panelists weigh in on what work, if any, they or their firms have done and if they think this might be a useful avenue to go down.

GREG AARON: One of the balancing acts that you have to take into account is if there's a problem, is the registrant responsible for it or not, because a lot of the domains that we'll see have been compromised in some fashion. Most phishing sites for example, exist on a compromised web server that a hacker has gotten into. And those domains therefore are not a good candidate for being suspended, because otherwise you might take an innocent registrar's domain offline. So you have to look at what is actually happening with these domains and make careful choices about how you mitigate the incident.

Now, in general mitigation should take place, I think, as efficiently as possible, but you have to be careful about how you do it. And that's why reporting out to registrars is important sometimes, because it's their registrant and they can help them get a problem cleaned up. Proativeness sometimes means looking for repeat offenders. There are certain people that as soon as we see them, we suspend their entire portfolio, because we know nothing good is going to come out of it.

And then find portfolios, as I've mentioned, that are starting to be used. And then you can make an evaluation whether you want to suspend the whole thing and take out a whole batch of domain names that haven't been used yet.

In general though, I would say mitigation, whenever it happens is positive. It does perhaps save some victims even if it comes a little later. You have to be careful about how you're doing it though. And every registry and its attorneys and so forth are going to have something to say. So you have to be careful, but using good data you

can get a lot of clues about what's going on, and what to watch out for, and what name servers are going to be a problem, and those kinds of things.

I'll say as far a zone update frequency, what I tend to see is somebody will register a domain name, they'll put it on a parking server, you'd like the default server at the registrar. And then they'll switch it right before they want to use it. Frequent updates are not usually an issue. We thought it would be an issue, especially when fast flux hosting was a big issue a few years ago, but fast flux has not turned out to be a terribly pervasive abuse of the system. And so fast updates of DNS, not one of the biggest indicators of what's going on, I'd say.

MICHAEL YOUNG: I think it's important to realize that proactive versus reactive is a subjective definition to some degree, because reactive I think could mean you don't do anything until someone issues you a court order. That's one definition of reactive. Another definition is how long it takes you to react to take a domain name down, even if you have a program to do that.

So what's really reactive or proactive? I think in today's lexicon proactive means that you've got an active abuse program in place. You've got tools and data feeds to support you in managing your abuses, and you're taking action when you have the confidence that you've identified something that's abusive. So I would define that as proactive in today's world.

Now that's a changing target. When you think about how abuse was a number of years ago, we started out with people registering garbage domains, numbers, strings, and they became very easy to identify very quickly. So the bad guys, for lack of a better term, started to register names that looked like they could be real domain names for a real purpose using a lot of the techniques that Greg was describing earlier in our presentation.

Now, we see a rise in compromised sites, because they're realizing it's even tough to hide behind legitimate looking domain names and they come down too quickly for their liking and for their investment. So now the bar as well will just look for server farms and complexes that aren't properly secured and will leverage their good domain name, their brand and hide underneath the hood, so to speak.

So I think every hole we block with a plug is they're going to make another hole and that's the nature of the game. So I think we can expect that indefinitely.

CHRIS DAVIS: I find this to be a really interesting conversation. So unlike my distinguished colleagues here at the table with me, I don't come from Afilias, which is a giant company in the space. The Foundation, when we started that we originally thought, okay, we've got a lot of experience in going to guys like Greg and saying "Hey, can we take down this domain because it's camatic control for this botnet and run different sinkholes and work with different working groups, you know, Mariposa was one of the big ones that we worked together to take down.

So when Paul Vixie, who I absolutely think is an amazing person came up and said, "No, I don't think there is any value after 30 minutes," or whatever it was, I was shocked. Because I think — maybe I don't really follow his line of questioning or his reason there — but I think that a domain that is currently being abused and has massive amounts of victims calling to this camatic control via that domain, when you suspend it if there was not value there or if it weren't fixing a problem then sinkholes wouldn't exist, people wouldn't analyze that data, we'd still have a big victim collection behind that.

So that being said though, the goal of the Foundation was that we wanted to go from instead of being reactive and saying, "Okay, let's go take down these domains." It's kind of like instead of treating the malaria we're trying to identify the mosquitoes that are carrying it. So let's go after the actual bad actors and try to make it harder for them to register new domains.

And it has to be a community effort. We need guys these guys. We need open sources guys like us. We need CoCAA using… I mean we need this sort of group effort to get this done. I think that there's just way too much work to do to after it. But I don't know — that's fine.

EBERHARD LISSE: Just a small remark. As having personal experience in the treatment of malaria, I think just going after the mosquitoes is underrated.

STEPHEN DEERHAKE: My other concern with the approaches being taken here is with the issue of false positives. You're pulling data from different data sources

with different polices/procedures as how one can appeal to them and get your domain off their evil doer list. And I speak from personal experience; this can be a difficult and time consuming process.

And I'm wondering, given that you guys are using these data sources, are you working actively with them, or making plans to do so, to try to get some uniformity into how potential false positives can be identified back to those data sources and prevent the mess of takedowns that shouldn't have happened on the registry or the registry's registrars standpoint and so on?

GREG AARON: Each data source has a different model that they use to define what they think needs to be dealt with. And understanding how they do that is actually really important. Now one of the things we do is we have multiple sources to kind of help triangulate issues. Now Spamhaus for example will tell you that they have an extremely low false positive rate. So when they list a domain they usually don't hear from the registrant, trying to get it back on. And they do have some verification processes and some automation that's going on behind the scenes, but you have to understand what that is.

And then you're ultimately going to have to make a decision about how you want to use the data. You're going to have to define what your thresholds are as an operator before you're ready to take some sort of an action. Or what you can do alternately is to say to say a registrar, "Here is a URL. Now, we're not recommending that you necessarily take down the domain, but you or your registrant need to take a look at this," because in general something good will usually come out of that.

This is a classic risk and benefit equation. Now, one of the interesting this is, that as I started to look at zone files we see perhaps a million domains a year being suspended. We very rarely hear though about when it goes wrong.

Now sometimes it goes wrong in a spectacular fashion. T.CO is a URL that Twitter uses and it was taken by mistake by the registrar. Somebody at the registrar did not follow their procedures. Those procedures have gotten lots of links taken down properly for phishing and so forth, and it usually works. But every once in a while somebody makes a mistake and you want to avoid that whenever possible. In this case it looks like the procedures weren't followed.

But what it also tells us is for most of the time the procedures, once they're put in place actually work really well. So it's interesting that false positives do happen. We want to avoid them, but for the most part it looks like people once they become educated and they've got good data coming in, do a pretty good job of making decisions.

MICHAEL YOUNG: And just from a future prospective, like any product we have a product roadmap. And not all our features are in there yet, but one of the things on my product roadmap is a feedback loop. So should a registry operator or registrar feel that something was a false positive they'd basically… It's like any application it you want to report an error. And that will feedback in a report to the data provider we got it from. So they will get regular reports to say, "Oh, we think this was a false positive," which will help inform them and hopefully help them improve their detection procedures.

STEPHEN DEERHAKE: Any questions from the floor?

PAUL VIXIE: Thank you, Doctor. This is Paul Vixie, ISC. Since my earlier comment was misunderstood, first by Greg and then by Chris, I want to make it clear. This is similar to what I said yesterday. We are in a long game; we're going to play it in rounds. I love the fact that everybody is now talking about abuse. Several of us have been ranting about abuse for 15 years. It's wonderful that you guys finally see that there's money to be made in doing this well, so you're all competing for how well you can do it — so clearly that's where I went wrong 15 years ago.

But I said 30 minutes, and what I meant is not that's the way we play the game today. I mean that best case, after you guys do the best you can as far as killing domains, if you give these people 30 minutes they can change their game in the next round so that they will make all the money that they need to make within 30 minutes from the time they buy the domain until the time you kill it.

That means they will stop buying portfolios. You won't be able to cluster analysis. They will not buy it until they're about to use it. They will find a very different structure for buying what they need, so that you can't get any kind of advance warning. And that's really it — if you can get it under 30 minutes, you can hurt them. And you have to hurt them. You have to make it uneconomical for them to proceed or they will proceed.

So what I'm asking for is not so much keep it under 30 minutes because that's the game we're in now. I'm saying that if in the next round you're more than 30 minutes, we will lose again.

GREG AARON: So Paul, I agree with 100%. One of the workflows that does exist in NameSentry for registry operators is if they're confident on the source — and we bring the data in real time — if they're confident on the source they can create a workflow that logs into a EPP client, to their registry and suspends the domain.

PAUL VIXIE: Thank you. I love that. Please push that hard on your customers. We don't have 24 hours. We don't have a chance for the next shift to think about this. This is a fulltime all the time problem.

EBERHARD LISSE: What's the time to live on a zone if you kill it in 30 minutes and the zone gets to live more and that will haunt us this work?

CHRIS DAVIS: I actually just wanted to comment. So yeah, I totally agree with Paul now that I understand what it was that he meant. What we try to do as a group of volunteers is if we can identify an actor and provide per transaction intelligence to a registrar — this is not necessarily at the ccTLD level — or to a hosing provider or to a dynamic DNS provider where they can say "This person is signing up for hosting, and do you

have any intelligence on this IP address, this email address, this piece of information?"

And we can reply back in less than a second and say, "Yeah, well that IP was used for this, and this email was used for that." We're not telling you don't allow the registration; we're just giving you intelligence that we happen to have. I'm not asking you to take the domain down, because I feel that that was a bit of a losing battle. So the direction we decided to go was let's provide you the information and you can make that decision based on your own risk policies and exactly.

STEPHEN DEERHAKE: Any other questions from the floor?

EBERHARD LISSE: Have you already got an API — Domain Secure has. Have you got one? For example, CoCCA would like to integrate this, but it needs to be negotiated, discussed, not only as far as the NDAs are concerned. Also the API in itself should not be an expensive undertaking, but making use of your service is of course a commercial issue, but will you make an API available for example for FRED and others?

MICHAEL YOUNG: We do have a RESTful — I say RESTful, because REST is a little indefinite, but a RESTful XML based API now. All the features that you saw on the web portal are accessible through the API, so you can basically rebrand the product on your own portal. Particularly I think we intended that to be useful for stakeholders or potential customers like CoCCA and also

registrars. Registrars often prefer to have a customer portal that they're controlling the experience from or they can provision data into. So that was our thinking in… Yeah, we do have the API.

EBERHARD LISSE: Because the more data entries we get into one point for the registrar to use the more likely they are to use it. If they have to go and look onto that website and that it takes human intervention. If it's just something going on I can scare my registrars into acting.

MICHAEL YOUNG: Yeah, I was particularly sensitive to this having built some of the original EPP registry systems. I've spent a lot years in the trenches building up and having to deal with a lot of difficult integration work. So I wanted to make sure that I eased that burden as much as possible for people who want to work with the product.

STEPHEN DEERHAKE: Alright, if there are no further questions from the floor. I just want to thank the panelists and thank Gillian for sitting in for Garth. And I hope you found this useful.

EBERHARD LISSE: Okay the next item… No, we're not ready yet for lunch, uh-uh. Mm-mm. The next item is the host presentation, Jacques Latour from .CIRA will talk about it. Don't go just yet. We'll break for lunch at 1:00 and come back at 2:00.

[break]

JACQUES LATOUR: Hello. Hello. Hello. My name Jacques Latour. I'm with CIRA. And I guess we're happy to have all of you here in Toronto. The only thing we couldn't control is the weather, but tomorrow's supposed to be nicer — sunny.

MALE: That's what they said yesterday.

JACQUES LATOUR: Pretty sure. So for the host presentation what I wanted to do is just give you an overview of what CIRA has been doing for the last year almost. We've been doing a lot of work internally to pretty much rebuild the entire infrastructure, so I just want to cover that. And then I'll talk about the new architecture that we put in place, the new registry that we put in place. I'll spend some time on DNSSEC because we're almost done the signing part, but we've got a couple of technical issues that some people in the room might be able to help. And then talk about IDN quickly.

So what we did in the last year is we designed a new network architecture and infrastructure for the network. We wanted to leverage ritualization, high availability and all that stuff. And then what we did is we built a new backbone, which is a 10 Gig platform that with fairly new equipment, with Palo Alto firewalls, F5, NetApp and all that. So it's

pretty high end enterprise type registry by architecture. So we deployed that and we're still in the process of migrating the existing infrastructure to that. But the today the new registry platform is operating on that.

So some of you who are looking at Palo Alto firewalls, they're a fairly advanced firewall; we like them. But you can talk to us, the IT guys in the back about some issues we had with them and the benefits that they provide us. So this is the new platform that we put in place. The second thing we did, the main reason for doing that platform is we wanted to implement a new security architecture.

So the key thing is to have zones, real zones, with real security policy between the zones. And to have the policy to protect all the assets that we have inside the registry. So the database we wanted them to be in a secure location so that we could control who has access to the data base. And this is just a partial picture of our security [arch] picture, but we have a zone where we have the signers. They're in a zone with very limited access.

We have public facing services, and in a separate zone we have our business application. So public service are allowed to talk to the registry. The business application is allowed to talk to the base. So what it does is it creates a lot of control for us internally to protect the assets internally. So that so that was really important; that was the vision that we had. We'd build infrastructure and then we'd put that on top of our network architecture.

And then what we did is, driving all of this is we wanted to redesign the registry to be a three tier architecture. So that was actually an 18

month project. We took the old registry and rewrote and ported a lot of the coded to WebLogic. The only thing I don't like is the licensing fee, but the product is pretty good.

So what it means with this three tier picture is that we've got the different zones of the application implemented inside each of the security zones on the new network infrastructure. I think that was one of our visions and we achieved that last June. Internally it means that our software development process, they're way more efficient. The other big impact is less downtime when we do software release and patch. With MailAware you just put the code in and it's available. In the past we had to have change windows of an hour or two to upgrade java code and stuff.

So basically it's more modular, more control on our part. We can do more with less. That means less hardware and we have actually more processing power than in the past. I think we use about half the hardware with this platform than we did with the old registry. On the downside, all of this increased the complexity of the solution that we have. So we needed to do more training internally to get people up to speed on all the different technology, the zoning, and all of that. In all it worked pretty well.

So some highlights. We used the F5 to terminate all the SSL connection and then have WebLogic Plus through on two nodes. So we have two physical servers with two managed nodes inside, so it's high availability infrastructure. And now we've got stateless beans — not coffee beans. JMS that was a big thing — JMS to do all of our time processing stuff, the housekeeper.

The big think is the database. In the past we had Zillion storage procedure and Oracle. We managed to get rid of those and put that at the abstract layer in the WebLogic platform. So it means now we're less dependent on the database. And I didn't attend the last Center Meeting, but I saw the slides from [.e] for migrating to Postgresq, so that's a very interesting alternative. Right now we're in the position that we have a new platform with much control and it's a much better environment. So that project we have worked for the past 18 months; finished in June.

And then right after that we started to work IDN. So we are working actively right now on doing French IDN. Those are the characters you can see. Obviously we had to do something special; we couldn't do the normal way. So we decided to do something called Administrative Bundling, and I'll talk about that. So IDN is now in the [OTNE] platform, registrars can connect and test and do some development there. And target date is for January timeframe, early 2013.

So the bundle, this is something we did a lot of consultation out there in Quebec and across Canada for French corrector IDN and what the community said is they want to do bundling. So if they own either part of a domain that's in a bundle that nobody else can register a domain. So it means that the bundle, they're sponsored by the same registrar and by the same registrant. And there's no way around that, so if you own one variant then you own all of them.

So we got that to work. We got some external help to figure out what the best way to this in EPP and all that and internally, but I won't cover that. The biggest issue we have was domain transfer and we'd just

trying to update. So because it's a bundle, on our side we've got to make sure that all the domains in that bundle are updated at the same time. But it's also more work on the registrar point of view to make sure that they follow the right process to manage the bundles. I guess we got all the technical stuff covered; we're now in the testing phase with EPP.

Okay, so DNSSEC. I'm not exactly sure when we started working on DNSSEC, but I'm pretty sure that November 12, we're going to have our zone signed by then. I guess we've been at it for about a year and a half now, maybe more, and it's been quite a journey. September 4 we had our Key Signing ceremony; we did that CIRA. We had a bunch of Canadian government departments that were invited to the session. We had a script and then we spent a good six hours programming the five HSMs with all the cryptographic stuff. It was kind of boring. It went well and that equipment is now in production.

Our DPS is online. You can go look at it. But basically we're going to do the ZSK 30 day rollover and KSK every year, and as far as we know in the lab everything works fine there. So it took a long time to do it and the reason we did it is we use a different approach to sign. I guess the biggest thing that we did with DNSSEC is risk adverse. There's a lot of lessons learned from all the registries that tried to implement or implemented DNSSEC, had different kinds of issues and we took those issues in account in our solution design.

So what we ended up doing is we have Dual Independence Signing Engines. So we sign the zone using two different signers and we compared the output to make sure that both of the zones are good.

And if it's all good then we sign, so I'll cover that in the next slide. The key thing is around the validation process. So we sign the two zones and then we do probably ten different types of validation to make sure that the output is good. And then if we see an issue with either signer, that's not behaving according to the spec, then we don't publish the zone file; we stop.

So the reason we did this is because it's really important to make sure that the zone file we publish is 100% good. So if we have two signers that generate a good zone file and we compare them and everything is good, then we're pretty sure that it's a good zone file we're going to publish.

So risk adverse, we looked at all the known issues that occurred out there. There were DNSSEC software issues, bugs with software. There was key management issue — that happened a couple of times — there's implementation issues infrastructure wise and also operational issues. So we put all of that; we did a bunch of workshops. We built a very detailed functional specification for what our DNSSEC solution is and it looks like that.

So basically if you look at this picture there are two parts. The top part is the signer at our production site, the signer set at the production site. And we have our backup online signer set at our backup site. It's not a backup signer, it's live. So we top we generate the zone file up here. We send the zone file to all the signers. They all sign the zone file. So in this case the little box 2.1 means we sign with OpenDNSSEC. Here we signed with Bind. Level 2 validation compares both zone files. We run a bunch of tests and if it's good, we publish.

And now you get the output of the four signer; they're all valid zone files that can be distributed anywhere on our secondary DNS server. So one key component here is that the signers at our production site, they're live and the backup site has live signers. So if we lose a primary facility then all the keys in the signature and everything is up to date live and we know and we're confident that it can be used at the backup site, so that was an important factor for us.

The other thing we did is we worked a lot with the OpenDNSSEC group/company to address all the issues we found with OpenDNSSEC. So right now we're running version 1.4.0. It's alpha. It works very well for what we need to do. We tested it out. Even though they said it's not for production, but we're going ahead with that version.

So that's a list of pretty much all the validation that we do. So before we send the zone file to the signers, we verify a bunch of stuff — basic stuff. We make sure the zone file didn't change by a certain percentage or an amount. We make sure it's valid.

Then we sign the zone file and the key things we do in there is we do LDNS verify zone. So we do a bunch of tests like that, to make sure also with a valid DNS. To make sure the signatures are good, to make sure the signatures are valid. Make sure the integrity, that either signer didn't forget to sign a domain or a domain didn't get dropped. This is the sum of all the risk adverse items we detected to put in our validation engine to make sure that we're going to publish a good zone. So if anything goes wrong we don't publish.

And then I guess maybe in the future we're going to… These are all software packages that we wrote that potentially we're going to make

available to the community as a package to implement this type of solution if you want.

Challenges, we've had a lot of them, and we still have issues. Every day we find a new bug with OpenDNSSEC or a new bug in our implementation, a bug in our validation process, the signatures, you name it. Talk to Jake in the back and Paul [Vouter], they can attest to all the different issues we've seen out there. Name it, we've seen it.

General the observation I want to talk is that although we use OpenDNSSEC and Bind to sign the zone file, they both sign the zone file differently. So the outcome is you can't do a diff just like that between the zone files. They don't behave the same way and they don't do it the same, so that made our validation very complex.

And today, just last week we had a bug. So when we do a key rollover and a domain got retired before and it was put back after. Bind would use the old signatures that it's got stored with the old key and OpenDNSSEC would use the new keys. And then when you go to do a diff to do the validation it doesn't jive, so that's an issue we need to deal with; that's the latest one.

But overall at least the validation engine that we wrote detected that kind of problem and didn't publish the zone. So we're pretty confident that the process is right, but we still need to do a lot of work to make sure that Bind and OpenDNSSEC actually, there's more… All the software products are there. That's about it.

EBERHARD LISSE: Thank you very much. How many domain names do you have in the zone?

JACQUES LATOUR: 1.9 million something, almost 2.

EBERHARD LISSE: Okay. Ah, we've got 2,705 of them are signed. Any questions? Roy what have you got to say about this? OpenDNSSEC are us.

ROY ARENDS: First of all Jacques, well done. I've seen the design of your system in Prague. And I was thoroughly impressed. You've basically… You take [visioning] to the next level. Basically the whole team I understand went from understanding how DNSSEC worked to actually building an implementation in production that need an understanding of both OpenDNSSEC and Bind in detail.

The result of that is you actually have found some bugs for the rest of the industry who are using both tools already, if you know what I mean. So I find it very good what you guys have done. And I also really like the design that you've deployed.

JACQUES LATOUR: Thanks. That's Canadian.

EBERHARD LISSE: Any other questions? In the back. You were just scratching yourself or what? Hang on.

MALE: Jacques — wow, that's loud. I was just wondering if you brought the picture of your [mooses] that you killed.

JACQUES LATOUR: It's in my bag.

MALE: People need to see that. It's pretty cool.

EBERHARD LISSE: Okay, there was a question there?

JACQUES LATOUR: Yeah, so today we have an issue with OpenDNSSEC and Bind, so I'm not sure if you're all here, but I'd like you guys to get together and figure out how to make it work so that it doesn't jeopardize our November 12 date that I committed to somebody one step up from me.

EBERHARD LISSE: It's cool that we can say we did it at Tech Day. Anyway, I think if there's no more questions we'll break for lunch now. We'll be back here at 2:00. João will start with a Bind 10 update and then he can maybe reflect over lunch about the error of their ways as far as that .CA's thing

is concerned. So I know it's not Bind 10 yet. And then Paul Vixie will talk about Rate Limitations and then we'll take what's on the Agenda.

[break]

EBERHARD LISSE: Alright, settle down; sit down. People, do you want me to call you all by name? Is that a yes then? Can you sit down please? Welcome to the afternoon session. We had a slight content issue, so we have modified the agenda slightly João is going to speak about Bind 10, and give us an update. And then Paul Vixie, who I don't see just yet, he will tell us about rate limiting. Take your time; we're not in a hurry.

JOAO DAMAS: Okay, I won't hurry. So I hope everyone had a good lunch. Some of you are still having lunch. I hope this doesn't affect the digestion. So as Eberhard said, I'll be talking a little bit about Bind 10 today. And then if you have any questions about where things are going we'll have some time to talk about that as well.

So I guess most of you are familiar with Bind, in particularly the current version — no, you have not, right — and the current version, which is Bind 9 and the long history that software has. The software development of Bind 9 was started in 1998; that was 14 years already.

So some years we were looking at this and how the environment has changed, and decided that it will be time to do a new version of Bind that will be better adapted to the current way things are running on the internet. So that's what we are calling Bind 10, and it's a work in

progress, but it's making quite good progress. I'll talk a bit about that later.

So what is Bind 10? Basically, as I said, it's the next version of Bind. So it has an authoritative DNS server, which is at level function wise of almost what Bind 9 does, so it's DNSSEC enabled. It has a completely new architecture. Bind 9 works well, but it's this huge one piece of software that does everything. For instance the recursive and the authoritative server are both in the same place and sometimes it's not clear to users how to control one or the other, how to select which functionality you want, because by default Bind 9 tries to be helpful and do everything it can.

So that's changing in Bind 10. We are separating functionality to make things easier for people. Maybe there is a little bit of additional integration, but at least it also saves some mistakes that were common. Part of this new architecture allows us to store data, the DNS data the zone data in many new, different ways. Bind 9 was basically an in-memory huge data base. It did have a small API to allow for additional SQL database access, but it was not really used by anyone.

In Bind 10 we are readdressing this whole approach by making our different data sources first class citizens from the beginning. And in the current Bind 10 that we are working on we already support SQL back-ends and the memory data sources. So the memory data sources are similar to what Bind 9 has, it works very similarly. It uses less memory; it is a little bit faster. It will get even better as we work more and more in it. But it also supports SQL and initially this is in the form of SQLite, so that to show how all this can be done in the new world of Bind 10.

Well naturally it works in master and slave mode, which means basically it has all the functionality you need to transfer zones from one to the other, including the server to server authentication provided by [DC]. So that's all there. It also supports dynamic DNS, so basically what you have come to expect of a complete authoritative certain name server is there.

Because it has more variety, in the future we will be adding extra data stores for it — specialized. For instance one of the ones that we already thinking about is something that will give you the kind of speed that you have come to expect of things like NSD. But at the same time make it part of the whole Bind ecosystem, not just a single specialized server. So you'd be able to select which ones you use, depending on your needs at any given time. So it gives you more control, more flexibility.

As part of the evolvement there are certain, let's call them byproducts, which have quite some interest for anyone who does any technical work in the world of DNS. Namely, there is a full implementation of DNS library C++. The software that composes Bind 10 is implemented in both C++ and Python.

We decided to use Python for the parts that interact more with the user or do complex separations that don't really need speed, and use C++ for the parts where performance matters most, so it's a hybrid. And so the libraries come with bindings for C++ and Python as well, and they are available for making your own development. Eventually as it says, Bind 10 will become a full replacement for Bind 9. I already talked about this, so I will skip that.

Where are we at? The 27th, maybe it was the 28th of September we released the first Alpha release. We have been working on this for about three years. And we've reached the point where we are now, going from development releases, where you could have kind of a sneak peek at the state of things, into actual releases. So as in any normal software evolvement process we will go through Alpha, Beta, and then finally release it.

So we are right now at the first Alpha release, which basically means we are pretty confident that this thing works. There are some tweaks that need to be done, some things that need to be added, but it is a good time to start testing it. It includes the complete authoritative server implementation that I mentioned before. It can be used in production. In fact we are running it in some of our servers.

I don't know how many of you are familiar with AS112 project, but that's basically an unknown Anycast service that syncs RFC1919 in other .ARPA lookups. And it's amazing how many queries there are on the internet for stuff that shouldn't be asked. For us it's actually nice because it provides us with a very good test deck of real actually internet traffic that we can observe.

So this thing has been running on Bind 10 RFC for the best part of four months now. It's receiving about 60,000 queries per second, every second, and it's holding itself up. It's a very good test of exposing the software to all the crud that's out there on the internet, not only the correctly formed, well intended queries. We are now going to in the next week or so put it in one of the name servers for ISC.org itself, as we gain more experience and more confidence in the working of this.

So basically if we can trust it to run our own domain, I think probably you could spend some time looking at it and feeding back comments that you may have, things that you'd like to see added, things that you'd like to see different. The Alpha 2 will be coming soon. If you want to join the Alpha program and have kind of a privileged access to the engineers working there there's a URL there at the bottom where you can sign up and do all this.

Alpha 2 will be coming in a couple of weeks and that's some stuff that we have pending from Alpha 1. There is one piece of software that we are working on that I think will make a difference for bigger tests, which is a Bind 9 to Bind 10 configuration conversion tool, so that you can basically drop the stuff and have it working without having to spend any time doing the manual conversion. That will come in Alpha 3, which we are scheduling for mid November, so about a month from now.

Then after that as Christmas present sort of thing, we'll have the Beta and we are targeting the final release, so the kind of 1.0 version of Bind 10 for January, 2013. Like any software building process, dates can be shifted a bit because you can do all the tests you want or imagine, but when things are actually in the field is when you find out where the real problems might be. And that's why we have this extended Alpha and Beta period. So that date can shift depending on what happens when this software is exposed to even more internet traffic than it is currently.

We know of several people who are testing this already in their labs and so far we haven't had a lot of reports of anything that's not working as it should, so it's looking good. At this time we will also begin work on the

recursive resolver. The goal here is to basically redo the whole recursive resolver to make it more complete.

The world has changed a lot. Perhaps you think that for TLDs — which is where people are right now in this room — recursive is not as important as authoritative. I think that's perhaps something that needs to be reconsidered, because after all the queries that you are getting in your servers are sent to you by recursive resolvers, so it's quite important that these recursive resolvers out there behave properly. So there should be some interest also from the authoritative side on what happens in the recursive side.

We are also trying to do some fundamental research around this recursive resolver. The world again, has changed a lot in DNS in these few years. There are people using the DNS in a lot of different ways, doing things that are now accepted that sometimes perhaps some time ago were referred to as [asyllogics]. Talking about SURBL, Semantics, how do you map geographical information in the DNS and so on.

Well, do you think that's a good thing or a bad thing? The fact is that this sort of usage is out there and if you want the internet to work well, you'd better be prepared to support them. And that's what we are going to do, that's our intention with the Bind 10 recursive resolver.

One thing we did find out through all these years is that if you think performance on the authoritative side is important, performance on the recursive side is ten times as important, because the recursive servers are facing not a few queriers, like the authoritative servers do; they are facing potentially millions of clients. Think of a big ISP with DSL access, and all the changes that have been occurring in the browser world for

instance. It used to be that when you were browsing the internet a few years ago with Firefox or NetScape it would issue a DNS query when you clicked on the link to go to the next set.

Things like Chrome don't behave like that at all. They do speculative look ahead. Whenever you load a page it starts issuing the DNS queries for every link that's on the page. So it can be prepared if the user has to click on any of them. So any page load, instead of being accompanied by a single DNS query, these days could be accompanied by a 100 or so, so you have to take account of that.

And then there is the fact of how some authoritative servers are using DNS for load balancing and that implies that they are using for instance very small TDLs, which affects the load on the servers. So performance for these servers is quite critical and we had to spend some time thinking outside the box. We don't want to do what has already been done. We want to think about different new ways of getting this stuff done.

And of top of that of course now everyone who is writing a new DNS server, authoritative or recursive, has to deal with DNSSEC — it's here to stay. Everyone is deploying it and that's a new fundamental difference from the world as it was 14 years ago. So we are busily working on these. The authoritative is coming up really soon now. If you want to do early testing on it, now is the best time to do it — not to be too late to the party, so to speak.

There is a dedicated website for Bind 10; it's called bind10.isc.org. You'll find everything there, including if you have any bugs to report, the status of your bug, so you can get feedback to see how we are paying

attention to you. And that's a bit of a summary of the whole project. The last thing I would like to say is, I would like to thank the sponsors that have made this project possible. This is made possible basically by a coalition of some of the people in this room, ccTlds in different forms. Quite a few of them contributed financially, some contributed engineering. I would encourage you attend the Bind 10 site and look at everyone there who has contributed throughout these three and a half years of the project and thank them all for their assistance. If anyone has questions or wants to know where things are going or perhaps not going I'd be more than happy to… So any questions? Your Johan?

EBERHARD LISSE: No, in the microphone. In the microphone so the remote participants can hear you. I think there is only one at the moment.

JOHAN: So you say that the authoritative server is almost done and possibly the release date is in early next year. Then you said the recursive server is next, but you didn't give any date for what state it is in and when it will be ready.

JOAO DAMAS: Right. The target goal is the end of 2013, so the end of next year. But as I said also, we need to think out of the box, do things really differently, so there is a certain uncertainty on the final date. But the goal is the end of 2013.

EBERHARD LISSE: What SQL engines do you support?

JOAO DAMAS: So right now, out of the box, you get Sequel Lite. We are developing towards support for a postscript SQL as well. It used to be that the world was if you wanted to do anything fast with SQL, you'd pick My Sequel as an operative engine, postscript SQL has a evolved a hell of a lot in the last few years and right now we think it's a better fit.

EBERHARD LISSE: From my own impression that is correct. I've done some stuff on a MySQL and a postscript SQL engine. It creates a little bit more complex — postscript SQL is much faster than MySQL.

No questions? We need some more questions because I don't see the next presenter in the room. Have you got his cell phone number?

JOAO DAMAS: I do have his cell phone number, yes.

EBERHARD LISSE: I see in the door there — no, it's not him. There is another question. Thank you so much for saving me.

MALE: So Joao, please speak up to status and path for the other various components of Bind 10, like signer stuff, like DHCPv4, DHCPv6, etc. I

mean there is a whole bunch of different components planned. Where are they?

JOAO DAMAS: Okay. Absolutely, so I was focusing on the DNS part, given the audience. It is true that as part of the work on Bind 10, Bind 10 is also working at the implementation of DHCP. So it's going to be a full set. Frequently people use DHCP in conjunction with DNS stuff, but mainly not at the level of a TLD, but certainly at the level of a local area network.

So we are working on that as well, and we will have at the end of this year a working DHCP server — not a very full featured one, but one that does understand and behave correctly on a network, so it can be used to provision. That's also completely new architecture, which loops in place to allow for special provisioning needs, like you see at like a cable provider. So that's all contemplated there.

As for things like the signer, what we are going to do there is basically take what we have in Bind 9, all these concepts of in-line signing and all the key management stuff; we are going to basically extract them from Bind 9 and make a common toolset for Bind 9 and Bind 10. So we would leveraging the code from one for the other — of course they are developed in different languages, but that's not really important — and provide a unified sort of toolset of experience, of administration knowledge that you need to have to have the system working for both.

And when we do that we'll also add some of the missing pieces that you don't have right now in Bind 9, like the actual key generation itself and all its support for policy.

MALE: Okay, excellent. So what about HSM support?

JOAO DAMAS: HSM support, Bind 10 is using Botan instead of OpenSSL. Well, [Paul Vartis] doesn't like it but that's it. So we will be adding support for HSMs, yes, as part of that support for DNSSEC. So, I think he wants to…

PAUL VARTIS: I'd like to explain. Paul Vartis of [iTed]. At iTed people are really closed on which crypto libraries are allowed to be used, and which are certified and tested and severely tested. So I can tell you right now that for [Well 7], you'll have either NSS Open SSL or [Ileg-i-crypt] and anything else will not fly. So Botan would be unfortunate and OpenDNSSEC has that same problem right now.

JOAO DAMAS: Okay. It's interesting that you choose OpenSSL and not Botan, but okay, given the issues…

PAUL VARTIS: That's history, and certification, and money, and other things.

[background conversation]

EBERHARD LISSE: There's Roy over there.

ROY ARENDS: Thank you. I'm going to respond to that. My name is Roy Arends. I work for Nominet. There's nothing wrong with Botan. Botan is really, really good; it does what it needs to do. Soft HSM is using it. It works really well. If you want to for instance look at Soft HSM and how it's implemented the PKCS#11 client size. You might find that it's fairly trivial to build the server side of the PKCS#11 library. In fact you might be able to marry those.

JOAO DAMAS: I don't know about you, but what experience with have with support for HSMs using OpenSSL is that when you talk about PKCS#11, you have then to add a last name to it, because there are many, many different incompatible versions. And the HSM vendors in particular, they all say the support PKCS#11, but they don't work with each other and it's a software maintenance nightmare, and even to get the drivers in the first place.

There are people out there for instance that claim that they have load balancing, but if you look strictly at the standard, that's not even allowed, so there's lots of stupid tricks out there, done in incompatible ways and I think as he said, Botan has a better approach to allowing you to keep things to a minimum.

ROY ARENDS: And you mentioned, I don't know about you, and we've done some work on PKCS#11 and OpenSSL libraries, so there's one thing that you need for instance, to get Bind 9 to work with an HSM and it's an OpenSSL PKCS#11 engine. And that actually doesn't work that well. This has nothing to do with Bind 9; this has nothing to do with PKCS#11. They're just different concepts that are being married into each other with this engine.

And so what you end up doing with for instance Bind 9 and OpenSSL, is for every individual signature you have open a session, which is costly. And the moment you try to work around that you're going to violate the PKCS#11 principle. So it is good to see that there is going to be some basic PKCS#11 straight support in Bind 10 in the future.

JOAO DAMAS: Yes, thank you, Al.

EBERHARD LISSE: And then after him, not one more.

THEO: Hi, Theo here. Just a question on the support model around Bind 10. Do you want to perhaps talk a little bit about that, your training schedules, priority support, and that kind of thing?

JOAO DAMAS: Support, well we'll be supported by ISC like Bind 9 is the actual kind of commercial terms of that are in flux right now. We are discussing that; we probably want to talk to you, Theo, while you are here. It's changing — the model is changing, but the main thing is that of course we are going to be supporting it. Offering different levels of support; it will be backed by ISC and its support engineers just like Bind 9 is.

MALE: Just a quick note, I missed the start of the presentation, so apologies if you've already addressed this. But in Bind 9 the only way to get HSM PKCS#11 to work is to recompile it using the PKCS#11 option, which is terrible if you want to support multiple versions — like you can't make one version that supports both Soft HSM and some hardware vendor. And I'm not sure if you're addressing this in Bind 10 or if you have addressed this, but that would be…

JOAO DAMAS: I'm not sure either.

EBERHARD LISSE: Why don't you download it and try it out?

MALE: Excellent.

EBERHARD LISSE: Alright, next is Paul Vixie. Do you want to have a box that you present from? He will speak about rate limiting.

PAUL VIXIE: I talked about rate limiting yesterday, now I'm going to explain it.

EBERHARD LISSE: He's now going to explain rate limiting to us.

[background conversation]

EBERHARD LISSE: Okay.

PAUL VIXIE: Okay, we're going to explain rate limiting.

EBERHARD LISSE: Take your time.

PAUL VIXIE: That was not true yesterday. So the heart of the lot of the operational security problems that we as DNS operators face is the lack of admission control on the internet. The source IP address of a given IP packet can be anything you want it to be. And it will most likely make it out of your laptop through the local gateway, through the internet core, over to some destination somewhere, which is likely to be one of your name servers.

If that source IP address is not right for the network it came from, let's imagine it comes from some DDoS for hire gang, who is being paid to dump an avalanche of traffic on someone, possibly an online gambling site. They would use the IP address of the victim, the online gambling site in a packet that they would send you, hoping that you would answer that packet, which you will do, because by the time you receive it you have no idea that it's fake.

The only person who can prove that it's fake is the first one to receive it, the far end ISP, the origin ISP would have the option of dropping that packet. If they looked at it and said, "Wait, that's not the IP address that goes with this customer or with network or whatever. I'm going to drop that, because I know." So by the time it gets it gets through the internet core and gets to the far end, it really could have come from where it purports to come from. So you are in no position to know that it's a fake source address.

And historically speaking that means you have to answer the question. Your answer, especially if it's a DNS secure answer is going to include some cryptography, which is going to be a big couple of 4,000 bit signatures is not unusual, especially on a negative answer. Some DNSSEC answers are bigger than others, and we have played the game down at the bar downstairs, not here, but at the other ICANN and IATF meetings of saying, "Well, I could imagine a query that would generate an even larger response. Let me try that."

That's fine if the good guys are doing that. It's when the bad guys are doing it that they then have the ability through a very small botnet, maybe a 10,000 node botnet. If each of those 10,000 sends say five

queries per second to a selection of name servers, where those name servers have been chosen in advance, because they support DNSSEC signed zones, like say mine, ISC.org or RIPE.net.

If it's known to support DNSSEC and it's known to be very well Anycast, very well connected, a lot of horsepower, they can send 50,000 very small questions to a selection of maybe a hundred different name servers. Send a small number of packets to each one of them and they will each answer that 60 byte query with a 3,000 byte multi-packet response using EDNS and IP fragmentation.

The online gambling site could easily receive 50 – 60 gigabits of traffic. They might only have a 10 gigabit connection, or they might have 100 gigabits, but it's not real likely that their transit provider has an extra 60 gigabits of headroom. So once you do this you're going to fill up that link. You will cause congestion. You will cause that online gambling site to not be able to get work done, even though none of you operating those name servers were necessarily harmed by this. And certainly the botnet was not harmed by this because it's only five packets per second per bot.

Again, this comes down to the fact that IP source addresses don't have to be right. They can be deliberately wrong and the packet will still get into the core and get to the far end and get delivered.

So ten years ago I wrote an Advisory for the Security and Stability Advisory Committee it's "SAC 004" — you can Google it. It is ten years old this month. But it is as true today as it was it was the day it was written. This is biggest problem on the internet and there's no economic model under which it could be solved. We're going to have to

live this way our whole lives and our children, too. There's never going to be admission control. That means you guys in your name servers are going to have to find some way to differentiate between a packet that's part of an attack and packet that is not.

We've got a way. Vernon Shriver and I talked about this for about a year and a half and finally came up with some code, it's actually all Vernon's code, but it was originally my idea. Where we look at sort of the signature that each flow makes. And by flow I mean we're gathering together all of the responses that we would send a particular remote end network about a particular resource record or a set of resource records.

We're not looking at queries; it's important to realize. We very early on abandoned the idea that we would be able to look at the queries to decide what was a flow. We have to look at the responses, because it's the fact that we are sending the same response to the same network many times per second that begins to inform us. Even though we can't tell that some of the source IP addresses are fake we know that no reasonable recursive name server would have a reason to ask the same question that many times per second. That's what makes it look like an attack.

Now when you're receiving an attack, you have the problem of well perhaps there are some real queries mixed in with the attack. Just because the victim is having their source IP address forged, doesn't mean they're not asking real DNS questions at the same time. So we very quickly dismissed the idea that if we decided that an IP address was bad we would drop everything from that IP address.

In particular somebody might have a reason to ask a lot of different questions. Let's say that they are a mail server, or they've got a mail server behind them and they really might have a reason to ask say, the VeriSign servers for .COM for a lot of MX records per second. But the point is they should not be asking for the same MX record over and over again in this same second, because they have a cache — we expect that they have a cache.

So I guess putting these things together into these flows, sometimes called buckets, and then giving each bucket an allowance, called a Token Bucket Scheme. So every time a new second opens up on the clock you give each flow a new credit allowance of how many tokens they can use. If they run out of tokens you stop answering them. So the problem there is somebody could still forge a stream of questions that they wanted to somehow starve the victim of the responses to.

So if you know the victim is about to ask a question for fubar.com MX, then you might forge a whole bunch of questions for fubar.org MX, sending all of those questions to VeriSign, imagining that VeriSign was running this code, and thus the real question would go unanswered. And that's a risk we didn't want to take. That would be what we would call collateral damage.

So our solution there is to not drop every question or every response that's part of that flow. Some of them are dropped. Some of them are turned into what we call Truncation Indicators, where we turn on the TC bit, that's one of the bits in the DNS header. It tells the requestor that they should… Essentially it tells them they should try again with TCP.

There's actually a much more complicated set of rules, but that's what happens.

Everybody has asked me, doesn't that create the problem of increasing the number of TCP queries at the server that runs the rate limiting? And it takes a minute to figure out why it doesn't. When you're answering with these TC=1 packets they're very short. They are the same size as the query, so it's not an amplification in order to do this. And of course we're dropping some, so if we attenuate the number of packets, but we do not amplify the size of the responses.

But if the victim who is receiving you TC=1 packets, if he's not currently asking a question and waiting for you to answer it, then it doesn't matter that you send him a TC=1, because he doesn't have a transaction that needs to be followed up with a TCP session. So pretty much automatically, no, we don't create a SYN flood problem with this.

This combination has worked. I know that Roy over there from Nominet has been running this on a couple of Bind servers for .UK. And Matt, are you still in the room? Matt raised his hand yesterday to say that Afilias is running this on the .ORG .INFO servers and it's working. And the way we know it's working is that there were a lot of attacks going on that week. That's why we put the patch out instead of being a little bit more careful with further testing and so forth.

I'm pleased to announce that unlike some other recent patches that we've put out in a hurry, this one did not cause anybody to dump core; we've had no problems from it and it's working everywhere it's been tried. We are eventually going to I think going to propose this to the IATF. And after either one year, or 16 years the IATF will produce

something which might or might not bear some resemblance to what we proposed.

Meanwhile it's in Bind. It is an open standard; we are encouraging other server implementers to please give it a shot. The logic is not that hard and the amount of extra memory that you have to keep in order to generate all these flows and try to remember them is very small. In fact Vernon told me today that we've been grossly misestimating this. So this is me, he's the math whiz — I'm the one who made this estimate and I was wrong.

I thought that we might use as much as a megabyte of RAM per 20,000 queries per second and so you might need an extra 5 megabytes of RAM in your heap if you were doing 100,000 qps, which is more or less the benchmark for these servers. And 5 megabytes is small even compared to my phone, so I was not worried about it. It turns out that 5 megabytes is actually far larger than what you could conceivably see, so we're expecting it's more like a megabyte per 100,000 queries.

So really, we haven't found a downside to it yet, other than it can be bypassed. It is possible to craft an attack that goes right through this. And if you know what the rate limit is, you can just keep yourself to sort one less than that number of packets per second, spread your attack across more servers, make sure that you're trying all of the different zones that each server is authoritative for. You can still get around this and we're working on more logic to deal with those cases.

But it's very important when you consider… Let's say you're a top level domain provider who's considering some logic that will deliberately not answer some queries. You have to carefully study the possibility that

you will ever not answer a query that was important. You're not going to throw this logic in if there's some chance, some good chance of collateral damage.

So as we think about how to solve for the next round — you know we're playing a long game here. Our side has finally come to the table and we're about to play the first round. We want to make sure that every time we improve it, we first do no harm. And finding a way to deal with one of these spread attacks that uses a lot of different domain names and does not do anything bad to the real victim, whose IP address is being forged, that's a little beyond us at the moment.

There's a mailing list for this and it's called Rate Limits at Red Bar, which is my personal domain. We will move all this to ISC at some point. You can contact me or Joao. If you want the URL for joining the mailing list where discussing either the implementation or the specification, if you want to know more about how to deploy it and how to operate it you could ask us or ask Roy or Matt. They've all got experience with that.

I'm trying to think what else I should say.

EBERHARD LISSE: Don't look at me. I just work here.

PAUL VIXIE: I know you just work here, but you might be the perfect test case. What else would you need to know before you would deploy this?

EBERHARD LISSE: I have not noticed any denial of service on my take. And I don't really do research into this, so I don't really know of what's happening. If my Anycast providers tell me that they see this, then I would be all for allowing them to use my data to analyze the situation. One of them is you, so if ISC wants to do it on our top level, you are more than welcome to do it, just let us know what you see.

I think it's a very good idea. I never knew it was even possible to take a site around the corner, so to say, through the back door. Fortunately our top levels are very… Well, Anycast — there is lots of horsepower, so I don't bother about it anymore. But a few years ago it was a serious problem. We had some issues that we couldn't really resolve any of our names anymore and then we had to go to some expense. The company responsible for it actually has since offered me $100 in compensation, and a letter of not doing it again.

PAUL VIXIE: I think that's cool. You should take the money and buy a beer.

EBERHARD LISSE: We asked for one U.S. dollar, but they probably couldn't get that little into their budget.

PAUL VIXIE: So you bring up three important points. First, these attacks are usually not harmful to the name servers themselves. They're don't need to be, the name servers are not the target. They're using the name server as a reflecting, amplifying DDoS projector.

So you can think of a very well provisioned top level domain with a lot of very beefy computers and very fat links and they're Anycasted around the world. You can think of this sort of as an orbital death ray that is up there in the sky and it can reach anybody. And anyone who wants to trigger it to fire the death ray towards someone else can do so. I don't know about you — I would not to live under a sky like that, but that's where we are. So VeriSign, I'm talkin' to you.

But your particular Anycast provider, that is to say ISC, has been running this logic on the SNS-PB complex for some time, because that's where ISC.org is also hosted, and so we had to have this for ourselves. We have not put this in for SNS COM yet, partly because it hasn't been abused and partly because we have a little bit more due diligence to do with the companies who pay us for name service than with the companies who are public benefit.

But in any case the last thing I would want to say about this is recursive name servers also quite usable for this type of reflective amplifying stuff. And unfortunately the last time we surveyed the complete 32 bit IPv4 address space we found 16 million of these open recursives. We found 16 million places that if we send a packet they will answer with a complete DNS response.

We would like there to be fewer. I know that a lot of those are Bind's fault because our default used to be the open for recursion. We changed that a few years ago. There was a lot of hue and cry, but we did change it so that we will by default only answer queries from on the same network. Never the less it is a long tail on that; there will be a very long period of time that these open recursives are willing to do this

type of reflection. Fortunately a lot of them are not DNSSEC aware, so they aren't as useful as a more modern TLD authority server.

But in any case the logic that we are using won't work on a recursive name server, because they can legitimately receive the same query from the same stub many times per second. Therefore we don't have good math that tells us how to bucketize, and tokenize, and credit, and debit, and penalize the flows in that case.

We did have this on a recursive name server that we were operating, which was the DNS changer replacement name servers that ISC was operating for the FBI for six months or so. We had to turn those on because we were being abused. And it wasn't hurting us, but we were certainly dumping whatever it was — not very much traffic, 75 megabits of reflected traffic was coming off our name servers and hitting somebody. I don't know who.

So we turned this on, because for DNS changer, these people were already victims of malware that had reconfigured their DNS to point to these name servers in the first place. I didn't feel like giving them poor service was really going to hurt them very much more. Also we were about to turn it off completely at which point they would go dark, so I figured better we answer some of your queries between now and the end of the court order than none of them.

But that's the only case that I would feel save turning response rate limiting on in a recursive name server. So if you're running a recursive name server, what I advise instead is put an ACL on it, make sure it's not open to the whole internet. Put some packet filtering at the edge of your network. If someone forges your address, your internal IP address

on packets coming from outside your network, drop them, otherwise they can make you DDoS yourself.

And if you do those two things you don't need any rate limiting on your recursive name server because you will only be dealing with internal clients who are probably not going to be trying to use you in this way. And we will continue researching the recursive case, just as we're going to continue researching the case of the widely spread authority attack to keep trying to find a way to work around us. Because I already know what the bad guys are going to do in round two and I don't feel ready for it yet.

So there's the much longer, and probably in my opinion, more boring version of rate limiting. Roy has a question

ROY ARENDS: Roy Arends, Nominet. You mentioned my name during your presentation. So here's some feedback on what we've done. So two of our name servers were heavily abused in a reflected amplification attack. And so we deployed these patches on these two name servers. Not on all of our name servers, only those two that were at that point abused. We had several [packet raises] come and ask if we'd deploy the patch.

What you could see is normally before the patch and before these attacks on average, it doesn't really matter which timeframe you take, it can be a second, it can be an hour — about 15 – 20% of the packets were repeats. So these were misconfigurations in the network somewhere. One of the cool studies that we did is the moment you

deploy this patch you can see the real clients. So the sources that are not spoofed migrate away to different authoritative name servers, because the DNS protocol takes care of itself.

If one of the name servers is not responding maybe due to this patch, there is always another name server that will respond. So my advice is if not all of your name servers are being abused and what we typically see in these attacks is that two or three or four of your name servers, probably the ones who respond fast to where the client is. If they are being abused put the patch on that, but don't immediately put it on all of them. Give the real clients a method to recover.

We had a session at the Center Tech meeting, where both Paul Vixie and myself were at and a few others. Antoine [Presudo] was there as well. At that point I pointed out and I don't really want to repeat it, I pointed out a way to circumvent the patch and… Is this being recorded, or is this being transcribed — this session?

Male EBERHARD LISSE: I think it's being recorded.

STEPHEN DEERHAKE: This is being recorded and transcribed.

ROY ARENDS: Okay. So I won't repeat what he… But we've since then found a few other ways around this and if you want to discuss in private the next round of attacks, we can certainly help you, and I think you will scared of what we found.

The last point that I wanted to make, you mentioned recursive name servers, and I don't mean in the sense of open recursive name servers out there — I mean in the sense that it's very hard to deploy this specific patch for recursive resolver, and I understand why. However, Google, which has something called OpenDNS; they are an open resolver basically. They have publicly documented their anti-abuse strategies and it has some interesting rate limiting aspects as well. So maybe if you're interested in deploying this on a recursive name server, they have well documented information on that. Thank you.

PAUL VIXIE: Thank you, Roy. With regard to the OpenDNS and the Google 8.8.8 DNS, I understand that it's in the business models of a lot of companies to deliberately run a completely open recursive DNS server. I can see some advantages to that myself in terms of the telemetry we would then be able to collect on how the world is behaving. But that's not necessary. It's not something that everybody has to do, like running a recursive name server.

So we're not putting a lot of priority on solving that problem, because clearly Google has a way around it. Google's particular way around it relies on having a 24 x 7 knock that is watching for things, so that they can put human hands on them as well. Most people are not going to do that for their recursive name servers. So as much as we probably will get around to solving that for our commercial customers, I don't know that I'll be trying to solve that for the whole world.

And yes, Roy, I always love it when you find ways to break stuff. So please tell us everything you find and we will work together on what we're going to do in round two.

EBERHARD LISSE: Alright, any more questions? Thank you very much. I don't know how I figure out my name servers, but eventually I will start reading the menus. Anyway, our next presenters will be Julie Hedlund and Patrik Fältström. As I said we have from the ccNSO Technical Working Group a brief mandate from Conserve [ccNSO] to sort of breach the digital divide, or breach the constituency divide and therefore we will have invited the Security and Stability Advisory Group to give us an update.

PATRIK FALTSTROM: Thank you very much for inviting us. We who are here, at least we here on the stage from SSAC, there's myself Patrik Fältström, Chair of SSAC. To my left, Jim Galvin Vice-Chair, and Julie Hedlund, which is one of the ICANN staff that is supporting us. When looking around in the room I have several other SSAC members in the room. Can the ones that are members of SSAC just raise their hand please? Okay and the rest of you can now see who has their hands up, so if you want to discuss the topics that we're going to go through briefly, you know who to capture in the coffee break, etc.

[background conversation]

PATRIK FALTSTROM: So Julie's working with getting the slides up, but the first couple of things I'm going to say might be sort of boring anyway for some people here in the room. So what we'll do her is that we will try to first give an overview of what we have done in SSAC, what we're doing for people that don't know who we are, and then immediately jump into a report that we released a little while ago about Dotless Domains, because we heard that from you constituency that was the most interesting topic.

The complete slide deck includes more material than what I will go through, specifically a little bit dive into two other reports I will not go through those slides. Next slide, please

So SAAC was formed in 2001 – 2002, so we have been around for a little bit more than ten years. We do guidance and try to help any party in the boarder ICANN community, not only ICANN Board, but also other parties who are interested in getting some advice. We have returned documents which have advice that is directed toward anyone that is using the internet or wants to use the internet, people that want to register domain names, so even coming domain name owners are ones that get some advice from things that we have written. Next slide please.

We are 38 members. They are appointed for three year terms and we are, as you can see on the slide, we're rotating about four to five each year. Next slide please.

In 2012 we have four internal work parties. Inside SSAC we have the membership committee that Jim Galvin, the Vice Chair is chairing. And then we have three active work parties at the moment, Registration Data Validation Work Party, Identifier Abuse Metrics, and Root Key

Rollover, where you can talk to other SSAC members involving what these work parties actually are dealing with. We're also participating in other committees and working groups, where maybe the most interesting one here is that we of course are hosting the DNSSEC sessions for example on Wednesday. Next Slide.

We are doing these kinds of briefings at meetings, not only at ICANN meetings. We are for example, having a session and Paul Vixie is the one that runs the session at the Internet Governance Forum in Azerbaijan in a few weeks. Next slide.

This year we have managed to be extremely productive. We have so far published seven different documents, and you can see on the screen what they are. "Advisory on Impacts of Content Blocking," which is a continuation of SAC50, which is a specific request from GAC. I think most of us would like the people like you to read that and come back with feedback on whether you think we managed to cover most topics. Next slide, please.

So now, Dotless Domains. We wrote a report SAC53 on those domains. Next slide please.

And the background is that the number of questions we got about whether a domain name without dots in it, that's what we call a dotless domain, whether that would actually work on the public internet. For example, you see the question here: If I resgister "dot BRAND" will I be able to use the label "BRAND" in a URL, on the webpage, in the web browser, in a name and address on the right-hand side of an @? And if I do, what will happen if I do that; will it have any secondary effects? So

these kinds of domain names that don't have dots in them we call them a dotless domain. Next please.

So the finding is that the resolution of dotless domains is not consistent and universal. If you look at different web browsers on the same operating system, if you look at the same computer, same software on different local area networks, if you use different DNS Stub Resolvers and if you use different limitations of email, you will get different results. Next slide, please.

The overall issue — that we of course guessed when we started this investigation — but I must say myself that I was in person, a little bit surprised how strong this was. There is an assumption out there that if it is the case that you have a dotless domain that is supposed to be a domain name and identifier that is to be used in local scope. If you just look at the DNS protocol,we all know about the search part, but we also discovered a multitude of other sort of assumptions and conclusions that are based upon the assumption that a dotless domain is local. Next slide.

So recommendations that we have are that just because dotless domains will not be universally reachable, just because of the reasons I just mentioned briefly, SSAC recommends strongly against the use of dotless domains. We also recommend that the use of DNS resource records, such as A, quad A, and MX in the apex of a top level domain be contractually prohibited where appropriate and strongly discouraged in all cases. Next slide, please.

So that was the content of the actual report that we released. So what then happened was that the board passed resolution that requests off

to and I read verbatim: "Consult with relevant communities regarding the implementation of SAC053 recommendations. Provide a briefing paper by September 31, 2012 detailing the technical policy and legal issues that may arise as a result of implementing SAC053 recommendations, listing the options, if any, for mitigating such issues." Next slide, please.

So what then happened after that board resolution is that the ICANN staff opened a public forum on the 24th of August, 2012 to request community input on the SSAC recommendation. So this comment period that has been open that closed on the 23rd of September with a reply period that is now extended to November 5, so replies are still possible to post. ICANN staff is running that consultation. We from SSAC just like many others I know are reviewing the various comments that are posted and there's actually quite a large number of them compared to some other open consultations.

So this is where we are and that was a brief explanation of dotless domains. And now I open up for questions.

STEPHEN BOTMEYER: Stephen [Botmeyer] from Ethnic, when you say that ICANN SSAC is reviewing the commands, how are they taking into account? Because in this case for instance there is absolutely no consensus in the commands, there is no community consensus on this matter. So how the commands will be used?

PATRIK FALTSTROM: Okay, that's actually a good question because ICANN staff is getting the comments. And because ICANN staff is asked by the board to try to come up with a suggestion on what to do, so ICANN staff is the one which all the questions are directed to. SSAC is reviewing the comments; we are looking at the comments and we have not yet made a decision whether we should file a reply to the comments within the reply period that closes on November 5.

If we are sending in the reply then that will be directed to staff just like all the other comments and replies in this open consultation. So we are part of this open consultation period just like everyone else that sent in comments, nothing more, nothing less. Was that an answer to your question?

That was quick. I think there is one more person that wants to say something.

ANDREW SULLIVAN: My name is Andrew Sullivan. Something that still isn't clear to me… I think dotless domains are stupid, and bad, and they don't work. But something that isn't clear to me is why this is a security or stability problem. I've read all the documents; I still don't understand that. It seems to me that if people want to do a stupid thing and they want to pay $187,000 for the privilege, we should let them bankrupt themselves — I mean go nuts.

PATRIK FALTSTROM: This actual also a good question. It's actually great — a session where we get good questions. SSAC in general when evaluating whether

something has a security and stability impact on the internet, on the area where are chartered to have a look at, we differ between two different security and stability situations. One, if it is the case that someone wants to do something and it only hurts themselves. And the second situation is when they want to do something and it might have impact on a third party.

And some of the findings that we have in this case include issues, as you can see them in the report, that it was actually part of the discussions in telling SSAC exactly what you just said, Andrew. But we came to the conclusion that in this case there is too high a risk that actually third parties are having issues, and not only the party that would like to put a dotless domain in the top level domain.

One of the reasons for that is of course that if it was the case that someone… One way of explaining it could be if is the case that a top level domain owner put a record in the zone and someone tries to use it and you only have the options that either it works it doesn't work, I would like to put that, personally, in the first category.

But if is the case that you put something in there and a third thing is happening, that someone is coming to a third party's website, or like in one of the findings we had, that if you say some operating sessions and default settings you end up in the local security realm, which means that it changes your settings regarding virus control in your local operating system. That is a typical impact on a third party, but is not an impact on only the ones that are the register for the domain name. So that is a secondary consequence that we are evaluating when making a decision whether we should file something or not.

Maybe someone else wants to add something?

ROBERT MARTIN: Hi, Patrik. This is Robert Martin from Packet Clearing House. I know some countries already use some of these forbidden records. I guess you've talked to some of them; I wouldn't know really, because I have been out of that area for a while, but this thing about getting into another security realm is not really… If it's a security thing maybe it's the wrong place to attack it is in the DNS. Maybe you should try to reach into the browser community or whatever, where they actually use the security realms.

PATRIK FALTSTROM: Yeah, there were two questions there. Paul do you want to…?

PAUL VIXIE: So it's known that there are some 15 to 20 ccTLDs that have A records at their apex right now. And so there's a certain ambiguity, where if you try to reach http://and this two letter code you might get a local resource and you might get a global resource. There's not a lot of ambiguity about these two letter codes; they mean what they mean. Now if on the other hand you go to each http://sales and you might get a global resource or you might get a local resource.

And either way it's going to be trusted the way that a local resource would be trusted. Then we're going to have a problem. And I think that the dotless domain report clearly disambiguates between the concerns

we have about ccTLDs doing this versus the new gTLDs doing this, or any gTLD doing this. And certainly .COM has never done this.

PATRIK FALTSTROM: To answer the second part of your question, you also asked "isn't this just an implementation issue?" And I would say, yes, you're absolutely correct. The problem is that kind of implementation algorithm is deployed on for example all Windows systems that are deployed in the world. So long as long as you manage to get a service pack actually installed on every one of them, then we can of course move on.

But on the other hand one should not joke about this because there is an important distinction between whether it is a protocol specification issue where we're changing the protocol or if it is the case that there is an implantation issue. But there are vendors that have chosen to for example use the fact that you're trying to use and identify that does not include a dot. In many cases that does not even reach the public unicast DNS tree as we know about it.

In some cases it will use a multicast DNS query with this the suffix .local and other camo things. So we have to think about… So unfortunately, or fortunate, I don't know really what term to use; people have started to innovate already with name spaces where you use names that do not have any dot.

If it was the case that everyone we knew that dotless term or a dotless token actually did hit the DNS then we would sort of only have the search part issue, which might be bad enough, plus the security realm issue that Paul was talking about. But it's actually much larger system.

This is part of when I said that I personally was a little bit surprised over the implications of using a dotless sort of token.

[ARTIS RETTA]: Paul, [Artis Retta]. I just briefly wanted to correct a view that I heard earlier that you could fix this in a browser or maybe in the OS. This is everywhere; every single command on every single computer worldwide assumes that a dotless domain is a local resource. You can't change that with a service pack.

STEPHEN BOTMEYER: Regarding the security problem. I don't get it, because today if type for instance 'fr', just 'fr' in my web browser I get redirected in some cases — it depends on the browser; it depends on the local network. And it is already a problem today, even if 'fr' has no A or quad A recall, which means that I don't see the link between the observation that one label domain names don't work reliably, predictably, etc. – an observation where everyone agrees with. And the suggestion to add yet another wall in the Applicant guidebook, yet another layer of bureaucracy in the ICANN process, yet another prohibition.

Because this one prohibiting A, quad A, or MX at the apex of the TLD does not solve the problem. The problem the already exists because there is no rule for what to do with a one label domain name. And this problem, this security problem, if it is a security problem, happens to every one label domain whether or not they have A, quad A or MX.

JIM GALVIN: First of all I think that you once again mentioned the earlier regarding the existing TLDs and the accessing ccTLDSs, so I think we already discussed that. The second part regarding Applicant Guidebook — our interpretation of the rules in the Applicant Guidebook is not that our recommendation implies a change of the rules that are already in there.

JAY DALEY: I want to second what Stephen just said and I also think it might be very.. I mean the problem is still there and I think it needs to be specified if it's completely legal or legal. And if it's completely legal, it needs to be on a protocol level and I don't think think ICANN is really able to take that decision without ITF intervening in some way.

Also because some of these TLDs today, they pose a security threat already on the current model. I mean if somebody uses initials of whatever, ccTLD has these records, and he expects to reach a local machine — maybe he does, maybe he doesn't. It's a security problem. If it needs to be addressed, it should maybe be done properly instead of in an ICANN document where some countries doesn't even care what ICANN says.

PATRIK FALTSTROM: I encourage everyone that either they agree or disagree with the SSAC report to file comments or to reply to comments, because once again it's ICANN staff that is running the open commentary and they are the ones that should listen to your recommendations. So I hope you have filed both of your comments.

JORG SCHWEIGER: I'm Jorg Schweiger, SIDN. I have a question of where this sort of where this sort of policy might lead to. Because I sort of agree with the previous speakers that when it's on a protocol level this is completely legal, right? And as ICANN added giving guidance on whether or not something is deployed correctly or not. So I wonder whether the next SSAC group work will be don't deploy DNSSEC because there are a bunch of rooters, or a bunch of firewalls down there that have implemented their own filters and they're not going to pass DNSSEC. I mean if people want to do stupid things in applications, they'll remain to do so, and are we going to write a report about it every time?

PATRIK FALTSTROM: Yeah, we hear what you are saying and the only thing I can say as a response to that is that also for DNSSEC deployment we have discussions in SSAC whether it is something that actually do have secondary consequences that are negative. And so far we have not found that.

Whether you agree with our conclusion or not is a separate thing. And this is one of the reasons why we in SSAC write our reports and then it's up to everyone to either agree or disagree with that report. And in this case staff of ICANN have made a choice of issuing a public comment period just to be able to listen to everyone on their view on the problem.

ANDREW SULLIVAN: Andrew Sullivan again. Oh, I didn't say before, I work for Dyn. This has always been possible, because I always was able to put — if I run

example.com I was always able to put the COM label in there and then put example.com in my search path. And then I'd type COM and I'd get there instead of the TLD. I'm wonder the extent to which the problem here is not that suddenly we have these people wanting you to use dotless domains, but instead that they want to use dotless domains and that's taking a large number of labels and putting them into the root.

So if the problem here is not in fact dotless domains, but the massive expansion of the root zone. And I'm wonder the extent to which the SSAC is going to feel comfortable drawing that conclusion and therefore recommending that perhaps the continued expansion of the root zone is not such a great idea.

JIM GALVIN: So Andrew just to repeat, to see that we understood your question. Is your question whether we have combined the discussion on dotless domains with a discussion on root scalability; that if it is the case that we do have dotless domains in that case the conclusion is that we'll see a faster growth of the size of the root zone?

ANDREW SULLIVAN: Well, another way to look at this is the security problem that you said before, because it's got the impacts on people — and so that's the reason — it's a third party effect. The basis for that is that as a matter of fact the search path facility has always been there, bad an idea as it ever was, and people used it. And from time to time somebody would come along and have a bright idea of putting some TLD into a zone and then they'd run into this search path problem. And the answer was

always "Well, don't do that. It's not that hard to avoid the labels in the root zone."

If the problem now is that there are thousands of labels in the root zone, you can't actually give the advice, "Oh, make sure that your thing doesn't conflict with this other thing." Effectively what you've got to do is tell people you can't use the search path or you can't use MDNS or you can't use any of these things. And MDNS is actually worse of course, because if you name your machine COM you're just going to get there. I mean there's all these problems, right?

So the effect of this is that if you have a small root zone the chances for collisions are in fact much smaller. If you've got thousands of labels in the root zone, it becomes a practical impossibility to avoid that collision in leaf zones and now you're just going to run into this all over the place. So if the security problem is that it's having these effects on other people, then the answer to that is well, we should treat the root zone more specially and make it small.

MALE: So let me respond to that by asking a question back to the audience to folks at-large here. I'm just trying to figure in here the different...

[background conversation]

EBERHARD LISSE: I don't want to interrupt the discussion, which I like very much at the moment. He will be next eventually.

JIM GALVIN: One of the things that's interesting, Patrik said earlier that what we have seen to date is a lot of innovation in name spaces, if you will. So this where we have this problem comes about because this is what people have done. We've created these realms, these local name spaces versus an internet name space and things in between, and applications have bought into this, and operating systems have bought into it. And one could argue that this is a good thing obviously; this is what you want from the internet.

And we're butting up against that directly right here and now. So the question that I would ask is — and someone else said here in the audience — I apologize for my recapture. One other point too, why not just let people do what they want to do to themselves? I mean why shouldn't we just let them have a dotless domain and if it doesn't work for them most of the time, who are we to care? It's their own problem.

And I think the thing that's important to keep in mind here, is we do have a responsibility to the overall security and stability and that's really the position that SSAC is coming from. The recommendation is fairly strong and fairly forceful because it has a technical foundation. Given the direction that people have taken in name spaces, this is where we are. We have to make that observation that doing this intrudes on that.

But it's worth noting that perhaps people should be allowed to experiment and create name spaces like this and see what happens, and see how it works. Maybe it's an opportunity for innovation and it's an opportunity for people to examine it and consider how they would

innovate and better solve this problem. Maybe name spaces need to change in some way.

So I think that SSAC's recommendation to not go forward with dotless domains is still the right thing. But the other thing which is excluded in our current recommendation is it's a blanket prohibition of that and that includes not allowing an exception case of reaching out through the RSTEP process and asking for an exception, because you expressly want to take advantage of this opportunity and expressly want to do this.

And you recognize fully what that means to the user community. That your particular user community will not get a uniform experience and you're prepared to deal with that. And that's what you want to do, so you allow an opportunity for the exception for people to do that. Does anyone have an opinion about that? Any comments, do you agree or disagree that would be a good idea or a bad idea.

EBERHARD LISSE: Sorry for the microphone.

JIM GALVIN: I heard them say. He's asking is that default in the gTLD Application Guidebook now? I believe that is what's documented now. I believe you're correct.

EBERHARD LISSE: Stop please for the microphone, so that the remote participants can hear.

MALE: Are you then suggesting to release a www domain, because I would love to have [AdWords] on that?

JIM GALVIN: That's a policy issue and I'm not going to take a position on that. Paul has a comment up here though.

PAUL VIXIE: Hello again, Paul Vixie, ISC. To Andrew, what I want to say is that the definition of DNS at the time it was done was in counterpoint to the old host.txt and they used terminology that was well understood at that time as defining hierarchical names. So I think that local host often works because most of us put it into our local recursive name servers. It does not mean that anyone expected it to. The idea is that a hierarchical name has at least one dot in it and that's what DNS is supposed to contain.

So the fact that you can ping a two letter country code today doesn't mean that the protocol supports it. The presentation layer of DNS was never very well structured, but this at least was specified. So for ICANN to take a position that it's fine for people to try this, would be for ICANN to step beyond DNS, and say that it wants to have some kind of roll in names that are outside the DNS. I don't think we want to do that. I believe that the SSAC report on this is consistent not only with the Applicant Guidebook as written and as agree to, but also consistent with the DNS specification, such as it is.

EBERHARD LISSE: Oh, there you are, Warren.

JIM GALVIN: Now, it's Warren's turn.

EBERHARD LISSE: But not for want of me.

WARREN KUMARI: Warren Kumari, Google and also SSAC, and also responding to Andrew. So yes what you are saying about the big expansion of the root being part of the problem, that's largely true, or somewhat true. But the thing that's important to remember here is much of the problem comes from the fact that the expansion includes lots of generic terms.

So for example I have a machine at home now, called Apple — I also have an Ubuntu and a Windows machine — I don't have a machine called .ST, which used to be one of the ccTLDs with a wildcard, and so there's just much more opportunity for confliction. And you know, I'm not involved with the Apple stuff; I'm just sort of a third party in this case and am sort of affected because of that.

EBERHARD LISSE: My own take on stupidity comes from my profession — teenage pregnancy — is that stupidity is generally not a crime.

PATRIK FALTSTROM: And with that, it's 3:30.

EBERHARD LISSE: Anyway, thank you very much. Just one more question. Where do you see opportunities to engage regularly? I think we really have to exchange email on this. I want council and our technical working group and I, we all want to engage more outside of our own ccNSO sphere. So I am actually a little bit worried that you got finished too quickly, but then the discussion was extremely interesting. So I'm looking forward to collaborating as we have discussed for the future meetings.

PATRIK FALTSTROM: Yes, to people in the room I can say that we from SSAC completely support the work that you're doing. And what we have been talking about is to coordinate the agendas and see whether we can increase the technical discussions here, also the first couple of days. So I think we will work together before the next ICANN meeting in Beijing to see what we can do together. Thank you.

EBERHARD LISSE: And now that I know your face, I don't have to send you threatening emails that you must be on time and things. Alright, thank you very much. Our next presenter is Michael O'Connell from co.za — .Africa or dot What? Domain Name services.

They have rewritten the registry system that runs [Coza] for all of you that know this. .za has got 35 domain names and Coza has got about 800,000 now and they basically had to rewrite their whole thing. The

coolest thing that they have done, they written something like what is called a Policy Engine; you can plug in policies for different registrars and so on. So he will explain to us how it works and go ahead.

MICHAEL O'CONNELL: These policy discussions keep us employed, huh? Yeah, so business effectively designs policy or defines policy, which keeps us motivated to earn our incomes. Policy is the language of business. So business defines our policy, which keeps us as technical staff, employed and entertained. Policy is always changing; it's dynamic with the times. So as the industry changes, so policy will adjust as we see with the dotless domains.

Business in general is not interested in hard technical implantations, so the technical implementation should therefore not hamper future policy development. So from a technical domain we need to implement policy efficiently, effectively and flexibly. To do so we've implemented a hierarchical procedural policy structure using dynamic libraries. You can generate your own libraries in any language you want, which has been stored within the system that you're implementing.

We followed the Unix paradigm of micro commands to build to an operating system of sorts. So each library has dozens of small commands checking availability of a domain from auction checks to clearance housing to all sorts of things. So the engine that we've designed then allows for on-the-fly reloading; we don't have to take the registry down in order to change policy. So if there is an adjustment, or a bug, or an error, we can just reload the registry engine.

The policy itself can be stored within the database or the file system, so it can be scaled quite efficiently. The capabilities of the policy engine itself maps the business rules directly, so it's a word for word mapping between your business definitions and your policy structure. We allow for policy variables within the XML, which just allow for an easier way to maintain the policy itself.

You can also design your own third party libraries and call them as you need to. At the moment the primary language is in Python, this being a scriptable language makes it the perfect suit for something that can be stored in plain text.

In the front of the policy itself we have numerous of number of XPath validations. This allows for Regex control, error handling, existence checking and so on and so forth. We also allow for a UI based policy merging between two policies, so managing and maintaining your policies between operational servers and your production servers becomes a lot easier to manage when you're merging the policies from test of production or vice versa.

The UI itself has contextual popup helpers, which extract data from the PyDock in the Python examples. So when you're actually putting your policy together it becomes easy to understand what each piece of code does. And as I said this is language agnostic, but at the moment it's only in Python.

So the structure of the policy itself, at the root of it is the variables for libraries and the XPath. Then it goes into the object definition, in this case it would be an EPP under domains, context, hosts, but that can be adapted for any type of object. The next is the pseudo events. These

events would be your EPP creates, your EPP infos, your checks, but you can also define your own events at this level, so that you can hand things around based on your timing at a later stage.

Within each event you have a series of activities and this is where the recursive magic of the policy engine comes into play. These line numbers, they are effectively policy code. In this case it would be Python calls to a library. And you can then have a series of child activities, which through the recursive natures goes back to the activities. And you have billion branching so you can split on checks, for instance if the transfer votes have been received or if the Sunrise outcome has been established. And finally if any errors have occurred you can roll back, or commit, or deal with a problem as it stands in the policy.

So the user interface itself, you can see the structure on the left hand side. You can see the object, domain, create. And you can run through a series of validations there, such as subordinate or delegated host checks, whether the domain exists, availability checks, if it's in a reserve list. We also can perform an account check to verify that there are sufficient funds within the registrar account. On the CoCCa side we reissue a name server check just to check any sort of intrigues in your name service that you provided when you create.

On the right hand side you will note the context popup, very similar to an Eclipse based setup. The PolicyExec library is defined in the policy parameters and is just a prefix that you can define there. And you can code raw Python into the activities itself, which makes it extremely

flexible if there's an issue or an error in Policy which hasn't been coded for in the libraries, which can then be incorporated at a later stage.

When merging two Policies the UI assists you — steps through both Policies and defines conflicts, inserts or nonexistent entries. You can than tick via whizzy-wig controls which side of the Policy you want to incorporate. When you click down you'll have a summary of all your entries, so you can review your work.

For the gTLD launch phase — this is just some pseudo-policy I threw together. What this is does is it checks data periods for Sunrise. So we've got billion branch there, so it's before the first of May. If it is then we issue Sunrise checks. We're got specific data engines there for — this is an example for .Africa that we'll then do a reservation for .Africa names and then subsequent to that will be international names.

We'll then issue trademark clearance, which will then be moderated by the clearing house. After that we'll issue our auction tokens to our applicant or whoever, parties interested in this. You can see at the billion branch fails to truth check will go into a Landrush check, which will then go through a reservation notifying applicant, get the auction on the go.

On co.za we plan on implementing a closed redemption period. The current EPP deletion cycle is ten days, which we've had a number of complaints that it's been too short, so we are now pushing that up by 20 days. This is just an example of how simple it has been to implement this on a policy level without changing any server code.

So pending deletion, which I believe is a five day suspension, we can then issue a timer — now you'll note that there's a pending closed redemption name there — that name is an event under the domain objects. So in 20 days time the server will pick up pending closed redemption timer and it will rerun through policy and rerun that code that you see below.

So it'll do a dependency check to see if any new dependencies have been created in the meantime. It will then reserve the name and then remove any blocking states and delete it finally. This process can be cancelled at any time during the 20 days, either via a renewal command, or a cancel of an action. So we have an extension to cancel a pending action. You'll then specify the name that you wish to cancel and that either will result in being debited the funds for the closed redemption or you can delete it as you need to.

Following the closed redemption, we're then looking at an open redemption. That's in trial at the moment. We're looking at about the 1st of March next year. So if there's no collection within the closed redemption period, we'll then issue a timer for open redemption. And as you see the same event structure; we're now issuing a new event name for the domain object for pending open redemption.

And then we can issue the token information for auction and determine whether a winner exists or not. And if nobody exists based on that billion check, we can then release the reservation and put the domain back into the wild.

And that's it, short and sweet.

EBERHARD LISSE: Too short.

MICHAEL O'CONNELL: It's never a short enough presentation. Any questions?

MALE: Not a question, but a congratulations. That's a clever idea. Thank you.

MICHAEL O'CONNELL: Thank you, very much.

EBERHARD LISSE: Well, as I said, too short. Thank you very much. That puts us a little bit ahead of our time. And our last presenter is supposed to be here at 4:00. Oh, there he is, excellent. He didn't answer his email, but he raised his hand, which is excellent. So, very well. Thank you very much for being short and sweet.

MICHAEL O'CONNELL: Absolutely.

EBERHARD LISSE: Didn't put me into any trouble. Mikey O'Connor is going to give us an update on the DSSA, yet another acronym. Again, as part of our attempt to reach over the constituency divide. I'm just trying to make a bit of small talk until he does his… Are we connected? There you go.

MIKEY O'CONNOR: It seems to be working. Okay, I don't know about you, but I kind of like the idea of standing up, so if you can hear me okay, I'll stand. It seems to be my fate to be presenting to groups of people where I am the last guy presenting between you and the bar. So I can go as fast or slow in this as you want. I think what I'll do is I'll go through it very quickly and then I'll tease you a little bit. And if you're interested, we'll go deeper. But I've gotten through this presentation in seven minutes. And given the subtle clues that I'm getting from this audience, [snzzzzzz]…

EBERHARD LISSE: You have a full hour.

MIKEY O'CONNOR: I have a full hour and we'll see how much of that I use. By way of background — I'm going to skip a slide — most of you know that DSSA stands for DNS Security and Stability Analysis Working Group. Anybody in the group? I know Warren's here or he was. Just stick your hands up. There are bunch of folks from the cc that are in the group.

This is a cross constituency group, so there are members form the GNSO, which the organization I come from. I'm the Co-Chair for the GNSO group. Jorg Schweiger is the Co-Chair for the ccNSO. You earlier heard Jim Galvin on the SSAC; he's the Co-Chair for the SSAC. Mark Kosters is the Co-Chair from the NRO; he's not in Toronto. And Olivier Crépin-Leblond is the Co-Chair from the ALAC.

And where we came from was a sort of eventful conversation in Brussels where the ICANN CEO Rod Beckstrom kind of got up in front of you all and said, "The sky is falling. And you're Chair along with the GNSO Chair and a bunch of other people." I said, "Excuse me? I'm not sure that's necessarily the right approach to this."

And out of that fairly lively conversation emerged the DSSA and we've been at it for a couple of years. I'm going to give you a pretty brief update on where we're at, but I've got tools to show you that you can have for free. And I have requests that you can help us out a lot, and we can spend as much time as it takes to get through that, and then yield to the beer after that.

[background conversation]

MIKEY O'CONNOR: Introduction to DNSSEC is after me? Oh, that poor guy. I pity that. Okay, so here's what we've done. It's been about two years that we've been at it. Interestingly enough, putting these cross constituency things together is kind of tricky. ICANN's not real experienced with that, so it took us a while just to learn enough about each other's cultures to understand the differences and learn how to work together. So I put that on this slide as a big thing that we did.

We also had some work to do in terms of clarifying what we were going to do. We had a bunch of methodology to build and all of that methodology is out on our website and is there for you to steal. I can testify that we think it's pretty good, but it's not done. It's just a start

and if we have time and you're interested we can dive in a little bit deeper and give you a taste of it. And I'd love to hear your reactions. This is being recorded, so I don't have to take notes, right?

MALE: Yes, it is.

MIKEY O'CONNOR: Good deal, okay. Where we're at since Prague, which is the last time that I think I saw you all, is that we are sort of in a slightly lower energy phase. We put out a report just before Prague and one of my heartfelt pleas to you all is that — that report's in public comment right now — it's just like the SSAC folks who were here a few minutes ago.

We went out for our initial public comment and we got precisely one comment and it was from an inventor who was sort of pitching his security gizmo; didn't really have a whole lot to do with what we'd been doing. And we weren't sure whether just a little bit confused or using the comment stream as a way to promote his product. Anyway we would love to hear from more people than precisely one. So that's my first sort of heartfelt plea, is some comments from the constituencies. I am kind of nudging along my colleagues in the GNSO and I'll kind of nudge you all, too.

The other thing that we realized is there was one little part of the report that wasn't nearly as easy as it looked. And I'm going to spend a few minutes on that, coming up. And that's another one of these things where we could use your help. So I'll save that for a minute.

And then we're going to sort of stay in take it easy more until the Beijing meeting. And then once that meeting's done, we're going to make some choices about how we proceed and I'll explain why that's going on in a minute. So that's why the “if needed” in the 'Still to Come' list.

This is a slide that I think you've seen before, but let me just pause on it. This is one of the things that is out on our website that you can steal. There's a pretty cool spreadsheet that sits behind this that basically lets you answer for you own organization, you don't have to share your answers with anybody, it's just an Excel spreadsheet, and in fact the DSSA is not asking you to share them with us. But it's a tool that we built for our work that you may find handy. And I'll give you the link to the website at the very end.

But let me just walk you through this picture, which essentially… And Jacques is here, so we've got to blame Jacques for this — this is something that Jacques and Rick came up with — the compound sentence approach to developing Risk scenarios. And so I'm just going to read the sentence that's on this slide. You read from left to right — it says: Well, an adversarial threat source or a non-adversarial threat source. And then underneath those it says sort of the dimensions along which you want to evaluate these things.

So let's take the adversarial threats first. An adversarial threat source has a range of capability, they have a range of intent and they have a range of targeting. So how capable are they? How intent are they? And how targeted on your organization are they? And all these scales are little dropdown menus in the worksheet. And what we found was that this was a very quick and easy way to develop a lot of risk

scenarios, so that you could evaluate them. And again that's sitting out on the website for free.

So we have a threat source, either adversarial or non-adversarial. A non-adversarial one would be like a big storm, or a flood, or some other act of God, whereas an adversarial threat source is an adversary, a government, or a hacking group, or who knows. You've identified your threat source — what's the context? What's going on in your environment?

One of the things is what's called in the methodology preexisting conditions. And it's late enough in the day for me that I'm sort of blanking out on what those could be, but it's okay because there's the list in the spreadsheet and you can go read them and learn from that.

There's also the preexisting condition of what security controls are already in place in your organization? And there's a giant list of those that's expandable if you want to expand them to include your own risks, security control, environment, and there are things that we missed. And then finally there are vulnerabilities that you identify in your organization. And there's again, a starter kit for you to pick from. But in your group that's working on this, you may want to add some more that are unique to your organization.

Okay, so now we have an adversary that's coming at you in an environment. They could initiate a threat event which could result in adverse impacts. So underneath the 'Could Initiate' is the question Well, how likely area they to actually do that? And it varies depending on all the stuff that's come before. For example a freak storm — I think

a good example is the tsunami in Japan. There's a very low likelihood of that. There was a threat event, clearly.

The next one — and again the tsunami in Japan's a good example — very low likelihood, but very high impact. And finally you sort of add all that up. There's arithmetic in the spread sheet that lets you arrive at a severity and range of impact for that risk scenario. And that's all summarized in this one spread sheet.

So this is a tool that we built. And I'll get to what we came up with in our first round, in a minute, but I just want to kind of spend a little time letting you know that it's out there and that you're welcome to take it. There's absolutely no requirement to send any information back to us.

If you would like to share information with us, we have a very elaborate protocol via which you can do that and be assured that your information will not be shared outside of a very small group of people that have signed nondisclosure agreements with you. So if you want to share confidential information, we've got a mechanism to do that. But this is just out in the wild for you to use.

Let me just take a minute… We took this methodology that took us a while to build — it's based on a preexisting methodology, but we tailored it a lot to fit the DNS ecosystem. And we ran very quickly through it to come up with some very broad risk scenarios that we identified in our first report. These aren't really done; these are just interesting ones that we want to go deeper into. And we arrayed them on the standard consultant two dimensional matrix. I think Dick Hart started it, but we've stayed with that.

So top to bottom is the strategic sort of view as opposed to the tactical view. And on the left is the sort of slow moving stuff and on the right is the really immediate stuff. And we came up with five broad topics that we want to look at. And I'm going to start at the bottom of the list because those are the ones that we in the technical community tend to be most familiar with.

We tend to think about things like inadvertent technical mishap brings down the root or a major TLD, and you all are familiar with this — this happens. And this scenario needs some exploring so that you can manage the risk.

Working up the list attacks exploiting vulnerabilities in the DNS bring down the root or a major TLD. It was kind of sad to see Paul walk out the door because I wanted to pitch this to him. He's never seen this slide deck, but there you go.

Working up, a widespread natural disaster brings down the root or a major TLD. We're starting to get out of the very edge kinds of things. We're starting to get into regional kinds of issues that might affect multiple organizations or even hundreds of organizations.

And then towards the top of the list are the ones that we in the technology community tend to either think less about or shy away from, and yet we in the DSSA think these are interesting topics to explore. The first is reductive forces like security, risk mitigation, control through rules, and so on splits the root.

And these last two are out of a fairly recent ISOC report that came out a year or so ago that talks about… You know, sometimes you talk about

Layer 8 or the God layer, or the political layer, but these are indeed, we think, real risks to the DNS and need to be explored in more detail. We the community are sort of feeling our way through how we're going to do all that.

And then finally, the last one and the most interesting one from my standpoint is gaps in policy, management or leadership splits the root. And this makes the political people really edgy, and I like making political people edgy, so that's one of our more interesting ones from my standpoint.

This is all sort of old news, but I've got a little more time, so I'm spending more time on these slides than you typically have heard in the past. This is a picture that you saw the last time, but it's gotten richer. In the last report this only had six things around the outside. It's now got ten and I got tired of trying to draw ten-sided figures in PowerPoint, so I turned it into a circle when I drew it this time.

And this is my next question for you all, and it's something that we really could use your help on and that is, we realized when we were writing the report that there are lots of different kinds of organizations that play a role in the security ecosystem, if you will. And we when we were first writing the report thought oh well, we'll just put all these people on this diagram and it will be easy. And in fact on one call, folks on the call said, "So, Mikey, why don't you just put all those organizations on the left side — why don't you just sprinkle them around that diagram and come back and show us what you came up with?"

And we quickly realized that this is really interesting. It's really complicated because — and pick any one of these. You all are the ccTLD registries, so you can put yourself in this picture if you want. But there are lots of other organizations. There's Oarc, who's one of the sponsors of this two day session, which by the way, I think has been fabulous. I've been sneaking in the back when I can, to sort of get a feel for what's been going on and I think this is a wonderful kind of rescission, I hope it continues.

But look at all those other ones. And so one of the things that we are doing in this sort of pause period is surveying people. Anybody who wants to participate at any level, either as an official person representing a constituency, or as an individual, or as a member of a corporation, or whatever — we're just asking where do you fit? Who does what on this picture?

Because what we found is that there isn't really a very good inventory of who does what in this kind of an environment, and it might be really helpful to know that. Not that we're going to be prescriptive. It's way outside of our remit to say who should do what, but rather just to find out who's doing what, much like what you all were doing all through these last two days, finding out about each other's products and services and organizational missions and all that.

We're just interested in who does what. And that's one of the things we're taking out to the community this time around. To do that, we built another spreadsheet. We're doing all this in easily accessible tools, so that people can steal them, modify them. You don't have to give it back if you don't want; you're welcome to extend it. But these are all in

Excel spreadsheets, they've been sort of tested in the nonproprietary framework and for the most part they worked fine. It anybody finds something wrong with it, let me know and I'll get that fixed.

And our goal is just to complete our report, but there is a sort of longer term goal that says these gaps, and overlap in policy, and so on is an issue that if we continue, we're likely to take up, and so this would be really useful information for us.

Now I've been saying "if we continue" and I need to explain why. In parallel with us the board has launched another working group, it's a subcommittee of the board, called the DNS Risk Management Framework Committee. And this picture describes the difference between what the DSSA is doing and what the board committee is doing.

The board is doing something that's broader than what we're doing. We are doing a risk assessment. We are assessing the risk, but we are not making any recommendations about what to do about the risks that we identify. We're just doing an assessment. And again this is all the way back to Brussels in that exciting conversation where the CEO was pretty excited about the risks to the DNS. We're trying to answer the question, what do those really look like?

The board is taking on a broader mission. They're going to lay out a framework that describes how that risk assessment is done and it's pretty likely that they'll steal a lot of our stuff. We're certainly hoping they will. If they don't steal it, we're really interested to see what their consultant is going to propose, because it'll likely be better. And that's great; we'll steal their stuff.

But they're also going to describe the mitigation, sort of risk planning layer. So once you've identified a whole bunch of risks, what are you going to do about it? Are you going to assume it? Are you going to ensure against it? Are you going to avoid it? Are you going to mitigate it? What are you going to do? That's not in our remit.

And then finally, they're going to take a look at sort of the monitoring part, because a risk assessment's only as good as the data on which it's built. And we're going through it the first time. Hopefully when the board gets done with their work, we'll have a process that can repeat forever. And so rather than put 50 people through a whole lot of work just to find it superseded, we decided we would wait until the board work got done.

So in one sense, we are a narrower scope; we are only doing risk assessment. In another sense we're a broader scope, because the board committee is primarily focused on ICANN, the corporation, whereas we are focused on the DNS ecosystem. And so we've been sort of doing the usual manage two projects in parallel dance. And I think we're doing fine, but what we wound up doing is describing what we're going to do in the context of that other project.

So across the bottom is the timeline for the board group, bringing up to Toronto. They have just selected their consultant, it's the Westlake group; some of you from New Zealand probably know those folks. Their charter is to describe that risk management framework between now and Beijing. And then after Beijing their charter is to launch this engine that they've described within ICANN, not in the whole ecosystem, but

just within ICANN the corporation. And their though is to carry on with that after Beijing.

On the top is what we are going to do. So what we've been doing is the gaps and overlap stuff that I talked about and this fairly thin soup right now that we've got in terms of public comments. Between now and Beijing we're going to refine our report based on hopefully some comments and hopefully some input in terms of who does what. And we're going to come back to you all for endorsement, hopefully before Beijing, we'll see.

In the middle is the stuff that we and the board group are going to do together. Right now we're in that aligning us and them to make sure that we don't waste each other's time or cycles. In the middle were going to be a conduit through which community input into that board work is going to happen. We're not the only one, but we're pretty well organized, so we'll be there to help with that.

And then to the extent that there is a community based piece of the ongoing process, we'll be around to help in any way we can to get that started. And then at the very end, we'll come back to that question, well, is it still something that we need to do — this broader risk assessment, or not. And so that's sort of our stake right now, is that we're sort of in the waiting — not waiting so much, as just much lower energy. We were a very high energy group, built a giant report, did a great job, but we're going to take it a little easy between now and Beijing.

So here's the last slide in the deck and I've got a way to go into more detail if you want, but let me just summarize sort of the pleas for help.

Again, we'd really love to hear comments on our report. We'd love it if you want to fill out one of our gaps and overlaps spreadsheets. And again, we'd love to hear comments on sort of our plan going forward as well. And there's a short URL to get to our page on the community Wiki. You really only have to write down the x/4AB5 at the end. So for those who are listening on the audio side it's the usual community.icann.org URL and then after that trailing slash is x/4AB5.

I think I'll stop here and let you all guide me. I'm seeing laptops get closed. I'm seeing beer signs in your eyes. So I will not be embarrassed or ashamed if you all say "That was great, Mikey. We'll see you later. Bye." But with that, maybe I'll just walk out in the middle of the room. I can do the microphone thing from there.

Oh except I… I keep forgetting there's some poor person behind me, right — the DNSSEC guy? No, I'm the last guy, okay. Any questions? I knew with that beer lead in that I'd get silence.

EBERHARD LISSE: I have a question. What's the difference between your group and the SSAC? I haven't really figured that one out yet.

MIKEY O'CONNOR: The SSAC — that's a great question. The SSAC is a part of this group. It's one of the five advisory committees that rolls up into the board. The SSAC, I think the easiest way for me to describe the difference is that we are a project. We have a beginning, a middle and an end, and then we drink beer. And I'm looking forward to that last, fourth phase. But we are not a permanent thing. We are a thing that has a start and an end.

We don't quite know when the end is, but there is an end. We are not an ongoing function.

The SSAC is an ongoing thing. It's forever. It's embodied in the structure of the organization, it's in the bylaws, it's got a very defined role that continues presumably until the end of time. Whereas we will be done; we are not structured as an ongoing thing.

EBERHARD LISSE: I mean content wise.

MIKEY O'CONNOR: Hmm?

EBERHARD LISSE: I mean what content wise, work wise?

MIKEY O'CONNOR: Well, we're doing a risk assessment. They are analyzing specific issues — actually I'm speaking for the SSAC, when there are SSAC members in the room who can do this better. And being a guy, I'm sort of making up what their mission is. I know our mission very clearly, which is we're just doing a risk assessment — one time through, what are the risks to DNS?

Whereas, I think, — and again, SSAC members, feel free to correct me — the SSAC analyzes specific questions that are put to them by the community. They could be risks, but they could be other technical issues that affect security and stability of the DNS. And their structure is

different than ours. They're not a cross constituency, per say. It's much more technically focused.

And they're also not very well structured to handle super confidential information from registries or registrars for example, which is part of the fabric of the DSSA that we haven't exercised yet, but we are built to handle. We have a whole protocol for handling confidential information and keeping contained within the group.

MALE: Can you elaborate a little bit on the confidential protocol stuff? As a group registry registrar I wanted to come to you with a GRIS question.

MIKEY O'CONNOR: Absolutely. If you have a specific question, we have a kind of short version of the protocol, which is… If you know Paul Vixie and you don't want us to know who you are, you can send it to Paul. Paul is prepared to anonymize your question and then scrub it and forward it to us. That's the sort of quick and dirty one, especially if you… And this is built into the report; we realize that some of the stuff in the report might raise a question in some of your minds that would be embarrassing to ask. And so we built that channel, so that people can come in to us without revealing who they are, and that's through Paul.

There's a much more elaborate protocol, which if you want me to, I could fumble around on my machine and bring up. But basically we've come up with sort of the standard consultants' two dimensional matrix that says there's sensitive information – let's see if I do this backwards; sensitive information, and not so sensitive information. There's private

and public and we treat each of those four quadrants separately. No, I'm sorry, it's always hard to do these without the slides.

There's sensitive information that's attributed. So attack data from a specific registry, that they don't want anybody else to know about except us. That's our highest and most contained. That goes into a very small subgroup and in each case the information provider gets to call the rules on how that's handled. So they get to decide who's in the subgroup. They get to decide everything about what goes on in that subgroup and they absolutely get to decide whether the information that's been synthesized by the subgroup is ready to leave the subgroup and go out into the broader group.

So the goal of the confidential information protocol is not transparency, it's protecting the data that is need for analysis, that' is being provided in that spirit. So that's the most protected. That group may choose to anonymize or synthesize that information into something that could be made public, and so they will prepare a draft of some sort. And again the information provider gets to make the final choice as to whether that publication is anonymous enough. And if they don't think so, then it doesn't leave the group.

Presuming it does, it's then public and it goes out to the public lists and to the DSSA, but it's not attributed at that point. That's the difference between those two boxes. This is attributed to the provider — this is not. Then it goes out to the world and the other quadrant is information from a provider that's attributed, but isn't sensitive and that can go directly to the world.

So basically what we're trying to do is build a wall for the information provider between the stuff they want to keep confidential, but share and that which can be shared with the rest of the world. And there's a fairly elaborate memo on rules and stuff that we've built to go with that.

And that is also available on the website. So if you have a situation where you, within your organization, are doing a project that relies on confidential information, doesn't have anything to do with the DSSA, but you want a protocol for handling that, you're welcome to steal that. We put a lot of work into it. We think it's pretty good, but if you want to use it for your own purposes, that memo and protocol is out on the site. You're welcome to take it and use it any way you want.

One of the things that we've been trying to do is build tools that you all can use in ways that don't have anything to do with the DSSA. I think that's it. Thanks.

EBERHARD LISSE: Thank you very much. So we are a little bit ahead of time. It's not going to be a big problem. Jay is going to give us his closing thoughts.

JAY DALEY: So I'm only going to be talking for about five minutes. So if you do leave now, I will think it's rude. So I'm going to talk about three things then. One is the last couple of days, but then I also want to talk about the value you get from this session and the future for these sessions as well.

So first of all, just over the last couple of days we've heard a lot about abuse mitigation, and abuse generally, and services, and research that people have done on abuse. And I think it's fairly clear to some of us at least that this is actually all about data. And that many of us are recognizing that as well as being service companies or service organizations, we're also data organizations. We hold a lot of data. If we all did full packet captures and kept all of our full packet captures all the time, we would be really big data companies quite soon.

And it's been interesting to have conversations with people about what they're doing about the data, how they're managing it, how they're storing, how they're capturing it, those sort of things. So ours is a relatively small registry; we have half a million names. We've just bought a large Hadoop cluster, half a petabye of disc space, terabyte RAM, that sort of thing for us to keep all our data on and explore an manage our data with.

And I know we're not alone in that respect, other people are doing things that are bigger and more detailed than that, and so it'll be interesting to see in a year's time, when we come back, when not only have we realized that we're data companies and that we need to treat data seriously. We've also begun to realize that there's gold in that data as well. And maybe our business models will have changed a bit around understanding how we can be getting money out of that data.

So that's my thoughts for the last couple of days, so now we'll move on a little bit to the value of this session. I'm a boy scout leader in my spare time. And when I need to get a vote out of a large number of children I have a very simple way of doing it — we all have a way. We

ask them to raise their hand and give us an answer between zero and five basically as to whether they like something. That means they didn't like it at all. That means they loved it, and any number of fingers in between, well you can guess what it means, so that sort of thing.

What we're interested in is the value of the session today and yesterday. It's expensive to travel. Some of you it's a couple of hours drive or five hours drive; for me it was almost 24 hours traveling to get here, that sort of thing. Business class, so not really that hardship, but still a long way to get here. So we want to make sure that this is valuable for those of you actually come this far and put this kind of effort into it.

For those of you have actually been paying attention for the last ten seconds, I'd love you to put your arms up and give me a vote between zero and five, and any number in between as to how valuable the session today and yesterday has been. Alright great, absolutely. Lots of fours, thank you, and fives, great. Thank you very much. That's very useful.

What we're going to do next time, because we're all getting into being data companies, next time we're actually going to run two sessions in AB Testing. So we'll do one in one way, one in a slightly different, then get you to vote, and we'll know which one we prefer. That was a joke. Don't worry. I know — technical people, we're difficult.

[background conversation]

JAY DALEY: Thank you. So the next question is about what brought you here. I'd like to know if there's anybody here who only came, or primarily came because we also held this in conjunction with Oarc yesterday? Can you put your hand up for us, so we can see? Okay, so that's maybe six in the room. That's very useful. So it's something we're going to try and work more with Oarc about to do this, but we find it particularly useful when we can do that as well.

So, you heard me mention, earlier when Patrik Fältström was in the room about SSAC, trying to create more of an agenda. This is something a bit wider that the working group is responsible for this meeting we've been working on. We're hoping to try to get the things that already take place that are technical, relabeled as being part of a technical stream, and then brought together as a relatively contiguous block in an ICANN meeting, so that basically we have…

If possible whenever we do the Oarc on a Sunday, the ccTech today on a Monday, then possibly Tuesday that fills in with such things as replacement for the WHOIS, or IDN registration things, some other sort of semi technical stuff like that with some DNSSEC stuff, then going through on to the Wednesday as well, and possibly broadening the theme.

And this is part of the general things that have been talked about making ICANN meetings more issues based, rather than constituency based, because I think that we're in a relatively privileged position within the ccNSO community that we get a day of technical things, when the gTLD people… You know, we need to feel sorry for them now,

don't really get a particular day set aside for technical things and so they're bits that spin around.

And so if we can bring something that will create something that brings all those people together that might be quite useful. We've started work on that and we're going to hope to try something much more — well actually have something up and running by Beijing around this. But again I'm interested to know from you how valuable you think that might be.

And again it's the same voting system. Okay. So if we are able to create a technical stream, whether you think that will be… How valuable do you think it would be for a start? That's one question; I'll ask you another question afterwards. So how valuable do you think that would be? Okay, great. Thank you.

So the second question, which don't worry, it doesn't require thinking about it, too much. One of the other things we're aiming to do is to create a space within ICANN through this technical stream that would enable more technical people to come, because they will be able to justify to their managers — because not all managers are as enlightened as me obviously.

Be able to say them, "Look I'd like to go to an ICANN Technical Conference bit for three days. Well I'm saying I'd like to go for five days to a very expensive country and I'll be doing stuff on one day, then I'll have one and half days, not sure what I'm doing. Then another bit of doing something and then you know, that sort of thing.

So I'd like to know from those of you who — just a simple show of hands — if you think this going to potentially make it easier for more technical people to attend if you think we have a more defined technical stream for the ICANN meetings. Okay, great. Thank you. Well, that validates things. Perhaps we should have asked you that before we started on this path, but nobody's perfect, so we're do the work anyway.

So hopefully then for Beijing we will have a slightly different approach to this. It will still work very much the same way as it has worked, but we will have integrated things a bit better with the other technical sessions going on, certainly in terms of the planning around the time, and the labeling of those other sessions.

So for example, if there is going to be a replacement for the WHOIS session, it would be nice to know there would be more than two or three ccTLDs in the room — that we would be able to run up as a whole mob in that room and set the agenda and help things go. And this session would have far more gTLD people in it than it currently does, because they would find value from doing it as well and we'd therefore get a broader set of presentations.

Okay, with that, just to thank the committee of Eberhard, Lewis, and Andre, and Don and Norm. And also to thank our host. Jack at the back, stand up, please. Our technical host for the week then. Thank you.

EBERHARD LISSE: Okay, that's it. Have a nice evening.

[End of Transcript]