Welcome everybody to the 37th ICANN meeting. As usual we had to wait for somebody called Rod to make his entrance. So there isn't much to say with regards to what is going on now. I think we have a very good agenda; we have a few remote presentations. I much more would have preferred to have those presenters be here but on the other hand it forced ICANN to work on the remote participation we wanted to have all along. So let's see if we can get that to work.

We have a nice agenda. As you can see, the morning meetings will start with Kim Davis and then Jay Daily will speak about security features and the application guidebook. Then by that time Louis will have finished writing his presentation and if Joe is late you will have to shift before him. Louis only came in this morning so his excuse.

We have a little bit of time, 15 minutes, if we go longer or quicker so we can adjust the time a little bit. But ICANN is paying for lunch. In the afternoon (1:48 – inaudible) and then Andre will tell us a little bit about some DNS attacks they had in Czechoslovakia, in the Czech Republic and what they did about it. Then as a late edition we have 1 ½ presentations from Chris Davis and somebody called Dagan at Georgia Tech. I only got the presentations this morning. They participated in the Mariposa.net take down and I think that is quite interesting for us because we are all affected by this.

Then we have a little panel about incidence response. Yuroito will talk a little bit about the DNSSEC and Roy Adams is the head of the Incidence Response Group and will sit on the panel and Patricia will remotely talk a little bit about what happened in Chile in particular with regards to the infrastructure and operations. Fortunately none of the staff there got injured or died.

Then Norm Ritchy will on remote make some closing remark. If he's not there or doesn't say enough then Phillip will say a few words.

So without any further adieu Kim will do his presentation. We must do it a little uncomfortably because it's the easiest way of getting it on the remote participation so they can see what is being said and in any case hear what is being said. Kim welcome.

Kim:

What I'm going to do is give you an overview of assigning the root zone. This isn't my presentation; it's actually a jointly compiled presentation by ICANN, US Department of Commerce and VeriSign. I am just here as the presenter and not as the author. So apologies if I don't speak with authority on some of this.

What has actually happened is whilst I'm responsible for managing the root zone at ICANN; we've had a dedicated team for assigning the root zone. I'll come to who is in that at the end of the presentation. But suffice it to say this dedicated team of full time employees have been working more or less non-stop on the root signing project. Obviously there is a nexus with what I

do as my job but I've been happily participating more as an observer than as an active contributor. I'm must monitoring how it's been going.

So this is a presentation jointly by those 3 organizations and also the design of the entire work flow of how the root is signed is a result of both cooperation between ICANN and VeriSign with the support of the US government.

I'm going to talk a bit firstly about the design. How did we design the work flow associated with signing the root zone? It is probably easiest to describe the design considerations by thinking about a few key words that went into the initial thinking and as the project got developed got flushed out a little.

The first key word is transparency. Signing the root zone is all about trust. DNSSEC is all about trust. To trust how the root zone is signed we need to be transparent about how it's being signed. So the processes, the procedures need to be open, as open as possible for the community to trust the signed root. And that really pervades the entire design of the root zone signing system we're deploying to try and be as transparent as possible so the community can have great insight into everything that is going on with signing the root zone.

The second key word is auditing. It is important to us that the processes we use are auditable. We want to meet industry standards, we want independent experts to be able to come in to test our processes, to view things like a key signing ceremony and say with confidence that this is a secure solution that meets the goals of signing the root zone.

Of course high security is essential. After all this is a security application. There is a variety of standards; there is niche standards and niche being US government agency. We've developed it to the high impact rating for system controls. The idea here is that we follow standards that are required for very high security applications so we can be confident that the system isn't compromised.

As you are probably aware from the start of my presentation there is a number of entities involved in signing the root zone, it's not just ICANN. I'll talk a little bit about what those roles and responsibilities are.

Firstly there is, of course, ICANN and ICANN here performs its role as manager of the root zone as one of the IANA functions. The role can be divided into a few different parts and this might explain why I'm not involved in all of that. Firstly, there is an entirely new part and something ICANN hasn't done before and it's exclusive to DNSSEC, which is managing the key signing key. This is sort of the master key that underlies the entire signing effort.

So managing that KSK is ICANN's role in this methodology and that's one of the new functions we've taken on. Secondly and this is more my area is the IANA will as it does today except changes to the root zone. And with DNSSEC the way the chain of trust is installed is using DS

records. So TLD operators that sign their zone will pass DS records to IANA for insertion into the root zone. That will be in addition to the current processes and is being rolled out as part of this project.

Obviously as part of managing the (9:01 – inaudible) in general we're there to verify and process requests to change the root zone. We're there to send requests to update the root zone to the Department of Commerce for authorization and then onto VeriSign for implementation.

Now the US Department of Commerce is responsible today for authorizing changes to the root zone. That certainly doesn't change with signing the root. So when we receive changes to the root zone whether it is inserting, changing or removing DS records, changing the KSK and doing other related DNSSEC updates, we will do that with the authorization of the US government. What the US government does is they act as; they authorize changes on the basis that they want to be confident that ICANN has followed its own processes for implementing DNSSEC. Again that is part of the trust underlying the whole process that we have published our methods and procedures and we intend to adhere to them in an open way and we hope that will be pretty evident. And NTIA should theoretically be able to approve all changes we submit because we'll follow a closely agreed set of processes.

VeriSign today is the root zone maintainer. They take changes that IANA sends and has been authorized by NTIA and implements them in the root zone itself. As a new role in the root zone management, they will be managing the zone signing key. The zone signing key is used for the day to day signing of the root zone. They will take any changes that NTIA authorizes as they do today and implement them in the root zone and then sign that root zone. They then distribute the signed root zone to the root server operators as they do today.

So I have a little process flow on the screen and it probably doesn't warrant going into any great detail. But the basic work flow is the TLD operators who have a signed root zone they transmit their DS records to ICANN and we then verify that and ensure it meets the procedural requirements. We transmit that to the US government for authorization and that's transmitted to VeriSign in an authorized state. VeriSign then generates an unsigned root. They use the ZSK they maintain that has been signed by ICANN's KSK and they then generate a signed root zone which is then distributed to the root servers.

So one of the key parts of the security of signing the root zone is protecting that KSK. After all, with a signed eco-system of zones all leading to the root zone that KSK is fundamentally the most important key to protect. Obviously, once the root zone is signed as TLD managers you'll be able to roll your keys relatively easily in the event of a compromise and transmit updated DS records to the root zone. But at the root zone we don't have that same easy possibility of rolling keys. So maintaining the security of the KSK is much more important.

There is a multi-layer security set up and I would be lying if I said I could step through every box on the slide and explain what each part means. But in essence we've devised a fairly elaborate

key control system where we've consulted with leading experts in the field; certainly VeriSign and their experience in managing private keys with PKO infrastructure have been a huge assistance there as well. We've engaged external security auditors that specialize in this field to design a system of physical security where we maintain the KSK in a very secure facility where we will perform key ceremonies that are very transparent. I think the slides will go into some detail on that.

One key part of managing DNSSEC for us is what we call the DPS, the DNSSEC practice statement. In essence, all the roles, responsibilities, the practices, how we do things is all written in a document. The idea is that as we maintain DNSSEC we will slavishly adhere to this document. This document can then be managed, it can be audited against and it's comparable to something that already exists in the security industry. With X509 certificates the kinds of certificates that are used for securing your bank transactions and so forth, they maintain what is called a CPS, a certification practice statement, on how they manage keys in much the same way.

One of the keys to us in managing that KSK is community trust. And in generating the KSK, in performing transactions on the KSK, we want to say with confidence that that was done in accordance with our practice statement, that there were independent witnesses that observed it being implemented and so forth. So one part of the proposal is what we're calling community trusted representatives. The idea is that we have community representatives involved in the events such as key generation, these community representatives actually by design have to be involved in the ceremonies. They will have keys themselves and these keys, they all have to come to some place and bring their keys with them and this is how they get access to the KSK.

So as staff we actually don't have KSK access by ourselves. We need these community representatives to come in and perform transactions with us. So this enables us to hopefully gain confidence from the community that this is being done in an open and transparent manner and there is no ability for ICANN staff or any other actor individually to do something unexpected.

I mentioned earlier that our aim here is to have 3rd party auditors to ensure we follow our practice statement to review everything we do and give independent advice to the community that we're doing what we say we do. Also, apart from those roles that have to actually be physically present to insert keys and so forth there is also an expectation that other external witnesses will be available to watch us do these procedures and report back to the community that yes I was there, I was a witness, I saw this whole process and I can confidently say the KSK was generated in the appropriate way.

DNSSEC protocol parameters, this is a bit more technical. The KSK is a 2048 bit RSA key and the idea is that it will be rolled every 2 to 5 years and that we use RSA 501 provisions for automatic key auditors and that the signatures be based on CHAR 256. The ZSK this is the key that VeriSign will maintain, 1024 bit RSA and gets rolled every quarter. So 4 times a year there will be some procedure where VeriSign generates a ZSK. They send it to ICANN and ICANN

goes through a process of signing that ZSK with a KSK and it gets transmitted back to VeriSign. That is done with those external parties. It is proposed that the signatures there also are CHAR 256.

Signature validity, DNSSEC key covering (17:34 – inaudible) validity of 15 days, other (inaudible) validity of 7 days,

Key ceremonies that I was mentioning earlier where we have access to that KSK will be generation of that KSK every 2 to 5 years and then once a quarter the signing of the ZSK for the forthcoming quarter.

So once you have a signed root zone and you have all these keys as an implementer of DNSSEC and as an end user you need a way of obtaining the trust anchors for the root. The idea is that ICANN will be publishing the trust anchors on a website both wrapped as plain DS records. We will sign that, we will issue a PKS 10 certificate signing request and basically that will allow other 3rd parties to sign the key as well and build trust, whether that's the KSK or not.

So deployment that is kind of…that's an overview of how it all is structured. In terms of deployment obviously the goal is to deploy a signed root zone. We want to deploy it with transparent processes, audited procedures and we need to have obviously deployment within validators, registries, registrars, operators and so forth. Signing the root zone is a critical piece of deploying DNSSEC but it's certainly not the only piece. We're not oblivious to that by any means.

Our goal is to communicate a lot and certainly we've communicated in many venues on our progress in DNSSEC, certainly a lot more than the standard domain community we normally communicate with.

So what are the anticipated issues with DNSSEC? The first one is having the DO Bit set to one. Those that follow some of the DNSSEC engineering manualists would have seen on and off discussions about this over time. In essence there is a bit in DNSSEC called DO, DNSSEC OK which signifies that the server accepts DNSSEC records. It is expected that when with DO Bit set there can be a significant proportion of DNSSEC queries being sent with that set, certainly I think its 70 or 80 percent but I might be wrong. That kind of level. Then a potentially significant population of clients might not be able to receive the large responses that are sent back associated with DNSSEC.

Another key issue is that once you've deployed DNSSEC and the trusted key is implemented in DNSSEC users, roll back is extremely difficult. If you can go from a signed zone to an unsigned zone it presents extremely difficult practical challenges. So basically unsigning the root zone is a very undesirable thing to have to do. So once we turn it and switch it on finally and establish production signed root zones it's not something you ever want to have to step back from.

So a lot of the design of the process has been very cautious approach to rolling out DNSSEC signing with various things and I think we'll get to them in a moment. We're very deliberate and slow so that we can spot problems early and avoid this roll back scenario at the end. Hopefully we can capture issues early on.

So in line with that we've taken a sort of staged deployment approach. Firstly, we've been deploying it incrementally across the root servers. We've been signing the root zone now for I think since late last year but we've just started as of January deploying it onto the root servers themselves. We initially started with the L root and we followed that with the A root and I believe in the last few days we've deployed it on M and I root servers and there are obviously a number remaining and J root will be last.

The goal with this transition now firstly with having some root servers signed is obviously to capture any issues with those signed responses but also by having some not signed at this point it enables the client population that does not understand DNSSEC that might be having some significant problems with DNSSEC to still be able to find root servers that don't have DNSSEC enabled.

This relies, of course, upon resolvers having some kind of round robin or some other mechanism that means they don't always query the same root server.

I think most key to the approach that we've taken is what we've called the deliberately unvalidatable root sign. What this is, is even though we're signing the root zone today and even though you can query the root zone today and get DNSSEC answers, you cannot validate against DNSSEC because we've deliberately blinded the signature that you're receiving from the root zone.

The reason we do this is so we can test responses to DNSSEC packets without having people relying on the root zone being signed. So if in the coming months we need to reevaluate, stop signing the root zone for whatever reason, we don't have a bunch of early adopters that have already started using DNSSEC. So this is really one of the key aspects to the approach we've taken to mitigate possible problems during deployment.

So this is basically what the DNSSEC key looks like right now. There is a key and internally we know what that key is and we can test it in terms of our testing procedures on a daily basis to make sure it is signed correctly. But the version that is in the root server itself you cannot validate against.

So the approaches to deploy conservatively prevent community validators from falling more or less repeating what I just said, so that we don't have a situation where we feel obliged to continue signing the root zone during this deployment phase.

Another aspect is measuring the root servers as we've deployed it. A focus of root server operators has been monitoring the traffic, monitoring the responses, observing any effects of this signed root zone. There is an ongoing dialogue with the operator communities to assess the impact of this process.

Testing is also key and we've been involved in a wide variety of testing. Obviously, there is still a lot of work to be done. I think we're gaining experience every day and that's going to continue.

I imagine most people in the room are TLD managers, so most interesting to you is how does this interact with TLD management? As I said earlier, the key to DNSSEC is if you have signed your TLD you want to have your DS keys listed in the root zone. So we're adopting an approach where the DS Key management will be done very much the same as say NS records are done today. You will communicate those with IANA, IANA will process those and IANA will submit them to NTIA for implantation in the root zone.

We expect to be able to accept DS records 1 to 2 months before final deployment of the root zones. So it is kind of May timeframe. To be honest, I won't be surprised if it's much sooner. I think internally we've more or less finished getting ready for that. Once we're ready, we'll likely just open it to the TLD community. But it will definitely be at least 1 to 2 months before the root is fully signed.

Communication, a website has been set up and strongly managed by the 3 entities. There are status updates, documentation, we have drafts of a lot of the procedures and so forth. There is an archive of all the presentations that we've given. Hopefully there is a lot of useful stuff there. There is also contact information to that root DNSSEC design team. As I said, I'm not part of that but they're very willing to give feedback about the process.

I think one of the areas of emphasis has been communicating with non-technical audiences. I think that is something we'll ramp up as we get closer to signed root zone. Emphasis on communicating with non-technical audiences as well as technical audiences.

With respect to technical audiences, obviously there is discussion in ITF forums, non-ITF forums like DNSSEC ORK and I know there was an RSA conference that almost the entire DNSSEC team was at in the last few days. General operator lists and so forth.

So the draft timeline has been spelled out to the community already but it's worth reemphasizing. We've started signing the root zone on December1st last year. Initially that was just internal to ICANN and VeriSign but we're doing it on a per zone basis. Starting in January we started incrementally rolling out that deliberately unvalidatable root zone to the root servers. It's intended by July that all the root servers will have that zone. Then on July 1st, all being well, we'll move from using this blinded key to having a full usable validatable root zone and root zone will be signed and fully deployed.

So here is the deployment status as of 24th of February. I apologize that it's a bit out of date because these slides had to be vetted in advance. We've published the requirements document, high level architecture, those DPS's the policy and practice statements have been published. We've published deployment documents and so forth. So there is a whole corpus of documents on that website I mentioned a few slides ago. It's at the bottom of this slide as well.

And we've also been documenting key ceremonies, KSK facility requirements, and I think from what I understand most of the documentation has been published if not in draft form I think there is no major significant documents yet to be published. But there is certainly a substantial amount of documentation already published.

We've been doing testing obviously; data collection with root server operators has been done to some extent. We've done several KSKR exchanges which means that we've taken VeriSign's ZSK and transmitted it to ICANN and ICANN has signed it and successfully transmitted it back. We've also tested the deliberately irresolvable root zone and some other testing.

I think at this stage M & I have been deployed already, so I think its 4 root servers have now been deployed. So that is kind of my rendition of a presentation that would have been given much better by someone that was actually doing the work. But hopefully I've given you a sense of the work that has gone into signing the root zone. It has been a pretty big effort by a bunch of people from the different organizations and I'd certainly be willing to hear your feedback on the process.

You know how we can communicate better, how we can try and address any residual concerns the community might have. Just to give you a sense of who is involved it is a variety of people from different organizations that there is a lot of expertise on that team and we have expertise from the DNSSEC community as well as expertise from the security community as well that have experience with PKI and so forth. So it's been a really good collaboration between all the parties.

I think…I'm obviously directly invested in the outcome but I'm somewhat an outside observer to the process within the organization. I was participating in key signing ceremony rehearsals in the last few weeks and it is really impressive the amount of thought and energy that has gone into deploying what we've got. I think over the coming months it will be a really good thing to see.

So thanks for your time and I'm happy to answer any questions you might have.

Chair:

Thank you very much and that was a very nice and interesting presentation. So could I ask a question from the Chair first? These community key witnesses, what happens if one of them loses their key?

Kim:

I'm not sure of the precise numbers but I believe it's a 5 of 7 scenarios. So there are 7 keys and we need 5 of them. There are multiple layers of, I won't use the right terminology so I'll try and avoid using terminology but there is a backup scenario where there is a second set of keys for a second facility on the other side of the country. There are multiple layers of redundancy there so that hopefully we won't be captured by any one or any couple of people.

Curtis:

I'm Curtis from (32:22 – company name) that runs the I root and we are signing the signed root. I had a question on the interaction between the TLD's and the Registry. I know it was discussed previously and was wondering what the outcome was on how does a TLD handle if they end up with a compromised key?

Kim:

The basic scenario is if a TLD operator needs to do a rapid rollover due to a compromised key, they would need to submit a new set of DS records for the root zone. ICANN provides a 24/7 emergency hotline for TLD operators. To date, it really hasn't been used but in concert with deploying DNSSEC we're going to go through another round of advertising that to TLD operators so they're aware.

One of the problems though is that by virtue of the structure of the root zone is you have 3 parties. You have ICANN as a corporation, you have NTIA as a US government agency and you have VeriSign. As ICANN we can say to you that we have a 24/7 number and what happens is you go to a call center which is outsourced and what they do is take your information and they will get an ICANN employee. They have all the available staffs' mobile numbers, home numbers and they will call until someone answers. That someone will look at it and work on it straight away. That is our undertaking.

However, we will do our best to make sure that NTIA gets the requests. We'll try and raise them, we'll contact VeriSign, we can't promise those other organizations will do the same but we'll do our best to try and do that. But our goal is to jump on that immediately but we can't promise on how quickly it will happen.

Chair:

Any more questions locally? Is there anyone on the remote who wants to ask a question? Not really.

Male:

I have a question, what happens when you get close to the very end of this and realize that stuff isn't going well? Has it been decided yet who has the authority to abort the roll out?

Kim:

I think the decision to finally go live at the end is a joint decision between VeriSign, NTIA and ICANN. We're in daily communication between the 3 parties and I think basically there will be a consensus between the organizations in the end. Certainly I think we would want to be conservative and the intention is to be conservative but there will be a risk analysis on whatever we choose towards the end and a final decision will be made.

Oh the question was is there a percentage of resolvers, is there a specific benchmark that we're looking for? I don't think there is. I don't think anyone quite knows exactly what it will look like if we're not ready to go. But I think we're satisfied that we have enough experts watching and we can engage enough people towards the end in discussion to get a good representative answer when the time comes.

Male:

I take it there is no automatic movement of information from ITAR into the root?

Kim:

No there is not.

Male:

When does ITAR get pulled?

Kim:

That's a good question. My intention is that when the root is signed, shortly thereafter, we will make some kind of announcement to the community and say as we've noted we intend, the ITAR was only meant to bridge the gap until the root is signed. The root is not signed and we'll engage the community to say is it appropriate now to shut it down. I'll just reconsult and then on the assumption that the community more or less agrees that's still the right approach then we'll do so with I don't know 3 or 6 months notice. I'm not quite sure how long.

The only concern there is it's been expressed to me in the last 6 months that there are elements of the community that want ITAR to be a permanent service. And I don't know if the rest of the community cares either way or if they're adamantly against it being retained. I just don't know. So I think clearly we have to wait until the root is signed before we take action on ITAR. I don't believe it's pressing to turn off ITAR right away but it is one less thing for us to maintain once the root is signed. So as staff, my personal preference would be for it to go away.

Chair:

Okay I see no more questions, so thank you very much Kim.

Our next presentation is Jay Daily and he is remote, so I have to load and run the file locally. Jay can you hear us?

Jay:

Yes I can. Can you hear me?

Chair:

Not very well, come a bit closer to the microphone. Okay it's getting better. Can you see the presentation?

Jay:

Yes I can thank you.

So this presentation is on security related proposals in the draft application guidebook version 3. We are promised a version 4 at some point but version 3 is all we have so far.

So what I will be covering is some background and context of this. Then the 2 individual components of this, the mitigating malicious conduct memoranda and the high security zones memoranda, often that is the high security TLD's and high security zone verification program as well. Then some brief words about the TLD ICANN working group that is currently underway.

Hopefully you're all aware of the draft application guidebook version 3. This is the current, well consultation finished on this a month ago but it is the guidebook that is given to applicants and new GTLD's explaining to them the requirements to be a new GTLD. It has within it 2 memoranda, one is called Mitigating Malicious Conduct, which is a list of compulsory requirements from any new GTLD and the next one is this rather unusual High Security Zone Verification Program, which is a floated idea for an optional program for new GTLD's only where they can end up with a form of certification by proving that they are at a specific level of security. So the definition of security is quite unusual here.

There is an ICANN working group underway at the moment and was set up quite recently at the same time as another one about zone file access where people are currently discussing what their views are about this. I'll try to give some indication of the work of that group so far but it hasn't really delivered anything yet.

So I'm sure most of you are aware that 2008 was the year of security and so was 2009 and so is 2010 and probably so will be 2011, 12, 13 and so on. We know that security is a problem when it gets onto the front page of newspapers and takes away from the marriage infidelities of film stars. And security is a very big business at the moment. It is something that politicians are actively concerned about.

We have quite common; quite commonly we hear about high profile attacks and all of this together puts ICANN in an awkward position. We have to remember that ICANN is in a central position of influence. It is being threatened by the ITU that the ITU could do certain things better, which we know it couldn't but it's being said. And it must be seen to do something. I think I should be a little fairer and there are lots of stakeholders saying to ICANN it must do something about this, some governmental and some non-governmental, some groups involved in this area and other things.

And for us CCTLD's whilst the job application guidebook doesn't refer to use directly, I don't think any of us can be in a bubble about this. This may apply to us some day and some of us may voluntarily choose to do it and some of us may be forced into it by our government. So it is something important for us to understand.

My analysis here comes from my work in the Registration Infrastructure Safety Group. This is a group made up of GTLD's, CCTLD's, Registrars and Security companies looking at sharing data and working on common consultations or common responses to consultations such as this. I'm not presenting the risk view on this that is well documented and published and very carefully worded. These are my own views but they're drawn from the work I did within RIS.

I think we missed a slide, can you go back one? No okay right then Registry Operators. First of all, this is the first element from the Malicious Conduct Working Group. So the Mitigating Malicious Conduct memoranda, just to remind you this is a set of requirements within the draft application guidebook that are compulsory. Every applicant will be expected to fit with these and this memoranda. It is quite well written and it breaks down the individual elements that people need to be aware of.

The first of these is (43:26 – inaudible) Registry Operators. We know that there are already many bad actors running Registrars and they are being slowly worked out and gotten through the system. But if we're going to end up with a couple of hundred new GTLD's we need to be very careful that this doesn't happen at the Registry level. So this vetting of those people who are applying to become GTLD's both for people and the companies and doing the bidding.

In my view, this is generally a good idea but there are some problems with the implementation recommended within the report. This is something you'll hear me say quite a lot, it's a good idea but it's just the way they're suggesting putting into action is problematic. For example, if you are just involved in a legal case then that can disqualify you from becoming a Registrar operator, which will allow somebody else then to just involve you in a legal case if they don't want you to become a Registry operator. It is making a judgment before the courts have done it, which doesn't seem appropriate.

The next thing is there is no mention of change control. We are well aware that Registries could in many cases turn out to be successful businesses or businesses that other people interested in ownership could change hands. It is at that point that it is quite important to ensure that the

vetting is conducted once again to ensure that somebody bad doesn't take over and this isn't discussed within the document.

Then there is no prevention on how to prevent gaming with multiple companies. We've already seen this with Registrars, people who set up shell companies and then shell companies within shell companies and other things. So we know people do it and we would have to put the same level of detection in place for Registries and that's not covered.

Our next requirement is that it requires DNSSEC and not just the GTLD must use DNSSEC but it must implement DNSSEC before or at launch. So when it launches it must have DNSSEC in place. Now this is a huge boost for DNSSEC and it's generally a good idea. But it is worth noting that current GTLD's and CCTLD's don't have this requirement on them, so we're asking something much more of new GTLD's and existing GTLD's.

And the Root Zone Scaling Study was quite clear that while there is no problem from having far more GTLD's or TLD's there is the possibility of doing too much at once, doing DNSSEC and IDN's and GTLD's will need phasing to ensure that these things work correctly. And by linking these 2 together the memorandum effectively prevents that type of phasing taking place.

The prohibition on wild carding, now the ICANN Board has already voted for this for existing TLD's, though there is uncertainty within the GTLD space as to how that will be implemented. Within the CCTLD space I believe the request has gone to the CCNSO Council and that has now pushed it out for consultation or is in the process of doing that. Within GTLD's there is differing views where the GTLD's, existing GTLD's simply cannot do it or whether they need to phase it out. We know that a number of them do it.

Now this is not a good idea but this is something quite important for people to understand how it came about. The Board vote was taken on a basis that came from the SSAC; it was not a standard community consultation process. So it didn't come from the GNSO community or CCNSO community. It came direct from SSAC to the Board and the Board voted on it and then pushed it down to the rest of us to then implement. Which many people might feel prevents the proper discussion that should take place about it.

Now in this particular issues about wild carding there are only a few people that disagree with this prohibition but if that same root is used for other things then it can be more problematic.

So as you're probably aware within .com at the moment there is a (47:53 – inaudible) by the Registrars and all the Registry sees are the name server records. Then when DNSSEC arrives they're see the DS records as well. And the DAG, the Draft Action Guidebook, mandates the use of (48:09 – inaudible). Now many people think it's a good idea, I certainly do, but we need to be honest about this. A (inaudible) gives better access to date but it doesn't give a better quality of data. The issue of security generally is the quality of data you get.

There are a couple of other issues within the proposal here that people have had an issue with. ICANN is saying as well as supporting who is, you will have to support any other protocol related that they say you must support. There are people who feel it would be unreasonable to just arbitrarily agree to support any protocol they're asked to do in the future.

Another point that many people from GTLD space make is that the who is, is not policed properly. We have regular situations where people are unable to get access to a particular Registrar who is, they deny access or the Registrar fails to provide a who is or a variety of other things they do wrong and very little takes place about that. As a result of that, there is a general feeling that if that were to be policed properly then much more benefit would come from that quite quickly than a longer term process.

The next one is central zone file access. This is a recommendation that all GTLD files be available through a single contract with a single point of contracting. From that then a security company can get access to everything. My view certainly is this is not a good idea. There is no diversity of the security of vetting procedures that need to take place within TLD's for people to have access.

I think we're all aware that very few of us CCTLD's allows zone file access at all and for good reason. Yet that discussion hasn't really taken place within the GTLD space in relation to this. There is a view that security companies need it, they say it's absolutely vital to their work and if you make it easier for them to do it by centralizing, then that will provide benefit. But as anybody who has done anything in this space knows before that if you centralize that access then it makes it easier for all of these companies that send out fake invoices to get hold of the zone files and make their life easier.

Now at the same time as the HSTLD working group set up by ICANN, a zone file access working group was set up. I haven't kept up with the developments there but they do have some draft recommendations out or a draft discussion document anyway that came out the other day.

Right abuse contacts and policies, now this is a slightly unstructured discussion within the draft application guidebook. There are 3 things they're suggesting there, the publication of abuse contacts within the Who Is database, mandated abuse policies (i.e.: a requirement be must have policies for handling abuse and some mandated content within them), and then publication of those abuse policies for Registries.

Now again my view is the contact side is a good idea. In fact, I was probably one of the people who suggested it to them. That makes it easy for people to find the abuse contact or Registry. But I have some problems with the rest of it.

First of all, what is so special about the abuse policies that they need to be put up there, rather than ordinary registration policies? There is no link on the Who Is database to, well these are the registration periods we offer domain names for or this is how you become a Registrar. Yet there

are views that abuse needs to go up there. I also think it is regarded as being within scope for ICANN to tell me what an abuse policy should contain. There are very different views in our community about how people should handle abuse. For example, some Registries are willing to be both judge, jury and executioner and take domain names down, others believe we should have an external process that makes the decision on whether or not to take a domain name down.

And my view is actually that those of us in the CCTLD community are far better at vetting these policies than ICANN can be in asking us what to do about them.

Okay expedited registry requests, this is very specific to GTLD's. This is where a GTLD Registry asks ICANN for contractual compliance relief. In other words, they want to be let off having to fill their contractual duties in order to mitigate a particular security issue. This is of course a good idea but unfortunately no details are provided on how this will happen or what threats will qualify. That potentially leads to some difficulties with one GTLD asking for something that others would had asked for had they known it was available.

High security zone, so this is the second memorandum. This is a voluntary program and people sign up to it and are externally assessed to show they have met the requirements. With that they then get a seal that they can put on their website saying we are a high security TLD and Registrars can put it up saying this TLD is high security as well. This only applies to GTLDs or existing GTLD to be able to do this.

It has a very wide scope of the things that it wishes to cover as well. So as well as doing the Registry specific IT and data security it also extends to general IT and data security in many ways. This whole document is a very messy document. What I am presenting next to you is a breakdown of this topic that is my breakdown and through colleagues within the Registration Infrastructure Safety Group but it's not presented in a document way. It is almost a discussion.

General IT security, now the document gives within it some tables of things it might be useful to certify against and from those I've extracted some key lines here to help you understand what they mean by general IT security. So security management, personnel security which goes into checking to see if you have a spy working for you, physical access control and then data collection, etc.

But as we all know there are already plenty of standards available for these. The most famous one being ISO-1799, which is quite established. There are already plenty of companies out there who are experts in auditing whether you comply with those standards. And there are plenty of training courses and books and other things you can buy to do that.

So this is really reinventing the wheel if we're going to expect ICANN to take some bits of these standards that it thinks is appropriate and put those into the requirements that meet the GTLD.

The next thing we have is Registry specific IT security, which includes things such as name resolution service management, DNSSEC deployment plan but the first question we have to ask is what existing Registries really agree on? I know from discussions with colleagues that I do name resolution one way. For example, I would only (56:26 – inaudible) and I know others who deliberately have (56:30 inaudible) with name servers. I think they're mad and they think I'm a lunatic.

So this is quite clearly there are lots of differences within these. The next question is what makes security so special from other operational practices of a Registry? For example, whether a Registry has a grace period or what type of billing process it does or what quality of data it expects. There are all sorts of things that could be looked at and used to objectively judge a Registry and yet those are not being considered. We're only looking at a few that might be useful for this wide notion of security.

The next one is Registry performance and it includes these marvelous gems of who is service availability, who is service performance leveling, who is service response time. The question has to be asked, what has this to do with security? Now if you go back to what security companies say they will say that access to the Who Is, is absolutely vital to them in order for them to track down registrant information. The claim they will make is that whilst the data given may not be correct, i.e.: it may be false data presented; it is very likely the same false data has been reused in lots of other places. There have been numerous cases where a particular line of address has been used in different Registries where it is clearly false.

But this is going really quite beyond that, me talking about who is service response time and various other things.

Okay verification of Registry and yes this is various occasion of identity for registrants of new GTLD. Now the views that other have expressed to me are, first of all, this is completely out of scope for ICANN. It would break the entire GTLD business model. The GTLD business model is not based on checking whether or not somebody is who they say they are. I know some CCTLD's do that but many don't.

There is also a very strong view that Identity fraud already works extensively within the domain name market. Many domain names probably the majority are bought with credit cards and people are regularly stealing credit cards and committing identity fraud in order to register a domain name with a stolen credit card.

So this is not necessarily going to tackle anything. Then we have the equal access requirement which is again very specific to GTLD's. This requires GTLD's to treat every Registrar equally. But if this comes in then it would enable GTLD's to say to Registrars well you're not doing the right kind of verification or you're not doing it at all so I won't deal with you. It begins to break that equal access requirement, which has been seen as a fundamental part of the GTLD market so far.

The verification of entitlement, now when I read this within the draft application guidebook, I think my jaw dropped open. Everything up until now has been about security and then snuck in there was this little line that I'll read. "Other considerations such as controls to address intellectual property concerns could be added as future components of consideration for the life cycle program." I certainly feel that this is not a basic security issue. I go further and say that I'm worried to see it actually included within what is actually a security related document. Intellectual concerns are entirely valid but there are processes for dealing with that. But to see it insinuating itself within a security document written by ICANN as a potential extension of a security based program is really quite worrying.

I have certainly seen people make comments that this is an indicator that ICANN may not really have understood just how much the intellectual property community can get into some of these areas.

Okay Registrant and Registrar interface, so this talks about setting a standard for the way the Registrants deal with the Registrars. For example, mandating too factual syndication or mandating out of banned communication such as leaving a text message telling you that somebody wants to change domain servers and you need to text back to authorize it.

Now this is another good idea but it is out of scope for ICANN. It is in my view something the Registrars should be doing and offering as part of their own differentiation. And Registrars have made the comment that by mandating anything this prevents Registrar differentiation. Now one of the problems that Registrars have had when tackling this so far is that all domains are not equal. Some people register a domain for a 2 month campaign, they register it for the minimum period that they can and are really now worried about it. Some people run a multi-billion dollar business off a domain and would protect that domain with everything they have.

Yet if we set a base line security standard we are creating all domains as equal. I think this is the view of many Registrars is that they should offer services to Registrants, allow Registrants to choose how they wish to do this.

Finally we get to the same problem of equal access requirements as to whether or not the Registry has to treat Registrars differently and some are willing to engage in this and some are not.

So a summary then, all together the discussion has been brought out from these 2 memorandums is a very good discussion. It is a very odd venue to see it taking place within the draft application guidebook. It is almost as though ICANN thought doing these things for existing TLD's is too hard. We'll just do it from the new GTLD's onwards and hopefully there will be enough of them that people won't worry about the old ones and we won't worry about CCTLD's.

But really there are things within that that need to be brought out and need to be discussed in a much wider forum because they affect all of us and yet they're buried away within this draft

application. There are some big issues within here as well. First of all, we have a change in ICANN's scope. I am very concerned that ICANN potentially is aiming to operationalize some changes that aren't in their scope which I would like to speak on by a diverse set of organization such as the DNSSEC and the high security zone.

There is in some ways a disregard for GNSO policy process within this. There are some things that are to be taken through the GNSO policy process but are simply being mandated within this. The next thing is a lack of empirical evidence. We do have many organizations pressuring ICANN to do things about this and that's very clear and ICANN has been very upfront who they are and what they've been saying. But we don't have the proper research that tells us whether any of these things will work. In fact, we don't have the proper research that tells us how they will work and what impact they will have on the market overall.

It could be that some of them, for example, if verification of registrant identity is agreed it stifles the new GTLD market all together and leaves the incumbent GTLD's running away with the market entirely. The other thing is the restricted scope as well. There is one slide that hasn't been brought up which is about orphan name services. One of the things that ICANN is recommending is the prohibition on orphan main server records, which they have mistaken for orphan (1:05:32 – inaudible) within a single registry.

So this is something where a Registry will be required to check that the main servers that appear within there for their delegations, if they're within their own zones must be for valid delegation. Now that is a genuine problem that does need to be tackled. But as our colleagues in Japan showed us there is an equal problem where people have created name service for delegation in other TLD's and those delegations do not exist in other TLD's. And yet there is no discussion about how that would be resolved throughout this document.

That form of restricted scope just listed new TLD's doesn't give us the real solutions we need there. But overall there is lots to think about. And finally just the ATSTLD's mailing list or working group, this is a working group with a lot of committed people in there so far. But it is struggling to find any agreements on what the nature of the solution might look like. Some people are very in favor of the certification program and some people would like to see a voluntary check list, some people think if we end up with a $20 registration for a secure one that is fine but some people think we'll end up with a $1,000 registration and that's fine.

So I don't see anything happening with that yet and the timetables on it look like it could be some time before it delivers. But for those of you who are patient and can read long emails I suggest you join that group and have a look at it.

So final slide please. So that's my presentation. I hope you all could hear me properly and excuse the coughing. Do you have any questions?

Chair:

Thank you very much and as usual a very interesting presentation. I'm also very happy that it actually worked with the remote presentation. Are there any questions? You have us all stunned it looks like.

Jay:

Or bored, yes.

Chair:

We are all in awe. All right thank you very much you can go to bed now.

Jay:

Okay thank you.

Chair:

As he's in New Zealand of course he was very late in the day, so he stayed up for us. All right and now we come to the usual host presentation. Our host Joe (last name) will speak about the set up and hope he'll also tell us a little bit about the technical situation. They changed the Registry software in the last few years twice and it's quite interesting to hear the experience that came with that.

Joe:

Thank you Chair. I'm going to give a presentation of ICANN Information Center, a brief overview of the ICANN Information Center set up, our current domains that we're offering, our marketing policy we are putting together right now and also talk about our Registry system of the GAC.

So ICANN Information is a multi-stakeholder registry, public and private partnership where we have members from different stakeholders, we have members from the telecommunication service providers, which is an organization of the ISP's and telecom companies. The Communication Commission of Kenya which is a regulator, Kenya Education Network representing academia, Domain Registrars Organization representing all the registrars, Kenya ACT Action Network, the Director of (inaudible) Government that is from the government and Kenya ACT Federation and the Kenya Internet Marketing Association among others.

Now like I said the Kenic Board is comprised of both public and private and most of them are from the government sector. However, the Chair and Vice-Chair are elected from the private sector just to balance the private and public sector. The constituent also allows for rotation of membership. For example, in my previous presentation those are the members of the Board right now. So after every 3 years we elect different Board members.

Now the objective of the .K registry, Kenic, is to operate and manage the K CCTLD, promote utilization of the .K name space, promote growth of the development in ICT's and the rural areas and low priced areas and also present Kenya in various international meetings and local events and conferences like the ICANN meeting, AFTLD, and the rest.

Now Kenic has been involved in the following initiative. Kenic was involved in the establishment of the East African Internet Governance forum and also the Kenya Chapter of the IGF. Kenic is pleased to host the 7th ICANN meeting. Recently last week from March 2nd to 7th Kenic was hosting the 4th AFTLD event where we held the AGM and IROC training. Kenic is also involved in the IP version 6 task force. We now have an IP version 6 test bed that we're going to roll out next month. We just received equipment from Cisco.

Kenic is also involved with the Regulator and also conducting right now research on CCTLD based practice. In East Africa CCTLD is Kenya, Uganda and Tanzania. Kenic is working with each of its partners, Kenya ATC Action Network on conducting this research. They have also facilitated the F root, J root and core main .net root servers in collaboration with the Kenya Internet Exchange Point.

Kenic also has a NTP server that most telco providers in the country use. Kenic is also the caretaker of the National Information Portal, Kenya.info. And Kenic was also very instrumental in coming up with their internet exchange point and they work a lot with the Kenya Exchange Point in building capacity in ICT in the country.

Now the model that Kenic uses is the Registry/Registrar model where Kenic has accredited Registrars, currently we have over 120 Registrars and then we have over 12,900 domains as of yesterday evening. We offer 2nd level domains not at the fast level. Now these are the various 2nd level domains that Kenic has been offering.

Since inception we've had 6 sub level domains Co.K for commercial, OR.K for non-profit organization, NE.K for infrastructure, AC.K for academia and SE.K also for academia, GO.K for government. GO.K is vetted so you have to be from a government institution for the government to get a GO.K. The same for the academia 2nd level domains, you must get an approval from the Minister of Verification to get an AC.K or a SE.K.

Now in January 2010 we migrated to the new system. Before we were using a Brazilian system and we have moved to Coca and after moving to Coca we have introduced 3 additional sub domains which is ME.K for personalized domains, Info.K for information, and Moby.K for mobile content. Now these 3 sub domains the main focus for these is on local content. This is the next phase we're looking at, especially after the landing of the 3 fiber optic cables in the country.

Kenic inception was in March 2003 after the delegation to Kenic and since inception when Kenic was started in 2003 we were using the Brazilian system. However, there were some drawbacks

when we were using the Brazilian system. We're not able to (1:16:05 – inaudible) to serve us better, in line with our policies. And Kenic right now is using the Coca framework since January 2010 when we migrated.

And some of the advantages we have realized after migrating to Coca are listed. It's easy to extend the Coca framework using EPP and since its open source we're able to extend to integrate with other payment solutions and other portals. Right now we are integrating Coca to some local payment gateways. One very common payment gateway is the mobile transfile called Empassa by Sampatico. So right now we are integrating with Empassa so that our Registrars and Registrants are able to pay for their domains using their mobile.

The Coca tool has helped us to have a system which is multiuser and multisystem so we can have Registrars have different accounts for their Registrants or Registrars can also go ahead and have resellers. So the system is able to provide that. The system is also flexible and we're able to introduce other 2nd level domains and we're also able to host other TLD's on the same system.

Presently, there is an enactment to the Communication Act in 2009 whereby the regulator is going to issue a license to anybody who wants to run a 2nd level domain. So Kenic we have already provided that platform using the Coca framework. Anybody who wants to run a 2nd level domain in Kenya will be able to run that. They will be licensed by the regulator and we can host them on our registry system.

The other issue we realized is the automation of course invoicing and domain clearance. With a BR system we had to flag the domains. For example, if one had to renew a domain we would have to go to the system and do it. Registrars didn't have a way of going to the system and flagging the letter for renewal. But now with the new Registry system, Registrars have the control of the domain. They renew the domain whenever they want to and if they want to suspend a domain for a certain period they can do that.

The new system is also prepaid. The previous system they had to register the domain and then pay for the domain later. Whenever they wanted to renew the domain they had to pay for the domain and then ask us to flag the domain. Now with the new Registry system it's a prepaid system as long as they have units on the account they are able to register as many domains as they can, they are able to renew as they can as long as they have units on the account.

The new system also allows us to use password tokens, which makes it more secure and also we have a better reporting system. So users, our Registrars and Registrants can be able to now use the system more flexible compared to the BR system we were using before.

To encourage uptake of .K domains Kenic has a discount policy for both bulk registration and if you're taking up a domain for longer years. Now if you're taking up a domain for bulk registrations between 0 and 50 domains, if a Registrar has between 0 and 50 domains they don't get any discount. If a Registrar has between 51 and 250 domains they get 5% discount, if they

have between 251 and 500 they get 7.5 all the way up to 50%. If a Registrar has more than 2000 domains then they get 50% discount.

We also have a discount policy for Registrants who take up a domain for longer. For one year there is no discount, for 2 years 5% discount, 3 years 15% discount all the way up to 5 years which is 25% discount. Our renewals and registrations are up to 5 years; however, with the new system we're able to go up to 10 years. But since our policy is up to 5 years we are currently doing up to 5 years. But the beauty of the new system is if we change our policy to 10 years than it's easy to integrate that.

The reason why we came up with a discount policy is because of the renewals of the domains. We realize that we were doing so many registrations in a month but at the same time we were losing so many domains. So we came up with a discount policy so that people can take up domains for longer years and they can renew their domains because if Registrars have more domains they're able to encourage their end users to renew their domain names.

Now Kenic as I said has introduced 3 new sub domains and we are currently doing a campaign to sensitize the community and the internet users the reason why they should take up .K domains. So our campaign is centered on .K and you have 3 main results why we're telling the internet users locally and also elsewhere to use the .K's. One reason you have a choice compared with the other GTLD's. You are saying that you are Kenyan and it's unique compared with the other GTLD's. I don't think that might be very clear. We have this brochure so you can have this on your way out.

The slogan for this is .K for every name in Kenya. Thank you very much.

Chair:

Unfortunately we have a little issue with the batteries for the handheld microphone. So we may have to wait a moment. Thank you very much. Can you tell us, and I can abuse the privilege of the Chair again, a little bit about how difficult it was to migrate from .BR to .K and how difficult it was to start from your previous manual system to get onto an electronic system in the first place?

Joe:

I think for us migrating from the BR system to the Coca was not a challenge because both of them are database driven. However, the BR system was in Portuguese so we had to sit down and learn some Portuguese and we also had to sit down and try to decode the database because even the database was in Portuguese. But after that everything was easy from there because we just migrate that data to the new database.

Coca we had to work with a local consultant because realized working with the team from Coca all the time was becoming a challenge, so we engaged one consultant locally who is also helping

the .NG, Nigeria also came up from the BR system. Our consultant has been very instrumental for us coming up with the Registry system.

After migrating to the new Registry system, of course, there are problems that come with a new system. So we first of all did a pilot with the Registrars. For 2 months we had a separate system running with the BR system. During that time, we were able to catch most of the bugs that were on the Coca Registry system that we learned.

Coming from the manual system in 2003 to the BR system, Kenic had to go around trying to benchmark with other Registries to come up with a registry system and we got good help from the Brazilian system, we got good help from the BR and that's how we came up with the BR system initially in 2003.

We are still fine tuning our Registry system. We still have a few issues to finish up especially integration with online payment. We are currently working with other local payment gateways to see where we can integrate credit and debit cards.

Any other questions from the floor?

Male:

Were there any policies in the Coca system that forced you to change the policies in the way you ran .KE?

Joe:

I will say yes and no because initially when we were setting the Coca there were some policies that were difficult to integrate our current policies. But we sat down with our consultant and Coca and we tailored…because the good thing the Coca platform is open source so we're able to change a few things and at least now everything is in line with our current policy.

Chair:

Any other questions? Thank you very much; it was a quite interesting presentation. I'll seek you out personally for 1 or 2 issues because I want to get hold of your data mining stuff. I'm quite interested if it's open source I might even be interested in sharing this with others. I've seen 2 things on your website which is quite cool which I would like to implement myself. Thank you very much.

The next presentation is from Louis Espinoza about the total cost of DNSSEC, the total cost of ownership and we must quickly get the logistics of.

Louis:

This is a topic that is a big concern for us. We are a small CCTLD.CR (Costa Rica) and we have around 12,000 domains right now and everything about DNSSEC is new for us and we're trying to define how we would cost that implementation for us.

I tried to establish 3 areas of investment or 3 areas of the deployment of DNSSEC on processes, technology and people. The point of view of process there is FC-4641 to establish some guidelines for best practice in this implementation of the DNSSEC and operations. The main things are the generation of the key process. One of the things that is important here is the calculation of the keys should be on a special device that has standard (1:30:21 – inaudible). After that the start of the keys is very important to be on the (1:30:26 – inaudible) device on something like HSM.

All of these things are expensive. The keys, the CS key and case key signing process suggests the signing process should be offline and then provide a secure copy of sign (1:30:56 – inaudible) to the master server. Most of the recommendations are based on the PKA certification management. The typical management of the keys in this kind of technology. Then those practices and processes is translated to DNSSEC and now the CCTLD managers must manage this kind of new stuff for us.

The point is trust. The security process is pointed to provide trust in the signer. Then how much can be invested in trust? I ask you how much authority invests in trust. How much a bank invests in trust? And how much a Registry invests to be a trust actor? That is one of our big questions.

This is our model of security in the field scenario. We only have, we are using (1:32:14 – inaudible) to manage the (inaudible). On the right side there is a model defined for the new deployment of DNSSEC. Now the second layer is not an option. Now the recommendation is that you have (1:32:44 – inaudible) server. In our case, we didn't really have it but now we need to set up that server.

After the server we have Fred again and we define to provide a new (1:32:59 – inaudible) to be used by Fred and then published to the server and then published to the DNSSEC. The signer is who has the keys and this is the most important thing in this process because it's at the bottom of the pyramid, the highest level of security. According to the recommendations this last layer of process should be offline even.

Then one each of these layers has a cost between deeper layer is higher cost. What about technology? Technology not always solves everything. In the case of technology to manage the keys is already in the markets, some HSM devices have security model devices around $20,000. The physical security to protect that HSM could be provided but could be expensive. It is a very high standard of security protection that may not fit the typical data center and to provide more security.

In the case of managing the table of DNSSEC with increasing data, it's possible to increase the memory of the servers and the capacity of the processors to manage the new sign. Typically the memory of server is low cost, not a big issue. The modification of software is relatively easy. It's a new field or something like that for the clients to provide the DS record and maybe a little modification in the process to capture that information. In our case, some part of the processes is running on Fred and this software is ready for that kind of fields, the field required to provide DNSSEC.

In the area of people the operations personnel have a cost of training and I think it's a relatively low cost. In technical personnel there are many new processes here and maybe the needs of new personnel to…this could be high for our organization. There is some change in personnel in our culture in this new development. This kind of change could be high in cost.

Some numbers, according to a study done on capital and operational expenses could be by example for a Registry between $250,000 Euros and 1.250 million Euros. That is a big number, especially for a small CCTLD. And for the Registrar they expect a 10,000 Euro cost in expenses and for the Record Operator between 15,000 Euros and 250,000 Euros. But my concern about this is not the Registrar or the Record Operator but it's about the Registry because it is what we run.

Then even the lowest expenses of 250,000 Euros is huge money for a small CCTLD and maybe for a medium CCTLD. That is only to provide DNSSEC. And that is needed but the value added is relatively simple for the user. So what is to be the investment in this development? What should be the value to provide a trust anchor for a small, medium or large CCTLD manager? What should be the level of security to protect the keys? Other recommendation to be followed or is needed more than that. That is part of my questions and I would like to know if any of the audience has an opinion or answers to these questions.

There are a lot of things about this scalability of cost. Some of the cost depends on how many domains you have. For example, the amount of memory, the amount of servers because the increase of the (1:39:13 – inaudible) to be high then you need more memory for the servers and more bandwidth. That cost is related with the size of the CCTLD. But some of the cost is no matter what the size of the CCTLD. For example, the management of the keys that cost is fixed because if you are a small CCTLD and you follow the best practice or the same procedures or process that if you are a big CCTLD and the keys are the same.

Maybe another cost that is fixed no matter what size of the CCTLD is the software modifications but this is not a big issue. The thing is the key management could be very expensive no matter what the size of the CCTLD it will cost the same. Then it is important to find alternatives to the numbers provided.

The idea is to provide some ideas and share some of your opinions on this. Like I said, the cost to be a trust anchor is very high for a small CCTLD according to the study. Then now I would like

to listen from the audience opinion about these questions I have. I don't have the answers right now because we are starting on this but I would like to hear something about it.

Ray:

Hi my name is Ray Adams and I work for the UK Registry. I've got a few remarks and hopefully a few answers to your questions. First of all, you mentioned RC4641 and that is actually a fairly old draft. Currently we're in the ITF rewriting that draft to become a better standard and more current.

Then I noticed the other thing, you mentioned the cost of managing keys, Nominate together with the Swedish and Dutch Registry have been involved in a project called Open DNSSEC and it's basically an open source tool that will manage the keys for you. It will talk to an HSM for you. So there is a very, very low start up cost to start deploying DNSSEC when you are using Open DNSSEC.

I also noticed 20,000 US dollars is what I understand for an HSM that's actually relatively high in the market. There are cheaper HSM's and for instance, if you want to use an HSM for key signing, which is what the DNSSEC is basically, for instance at Nominate we use a 800 UK pounds HSM called an SEA6000. It is an off the shelf hardware, it is certified and basically satisfies all the requirements that we had at Nominate. Now, of course, our requirements might be different then from other Registries. For instance, I know that VeriSign as a Registry they have a slightly different standard to adhere to. But I think what we have built will satisfy in general a lot of the requirements that a lot of CCTLD's have.

One final point, the latest version 970 can do automatic online signing. I know it's a slightly off the standard for 641 that says don't do online signing but if you're willing to go around that, BI-97 does actually do online signing. Thank you.

Louis:

Thank you. Any other comments? Okay thank you for the information and the reference of the RFC as I know is a new standard. The thing is we are really trying to use DNSSEC for matching the keys but the thing is I believe the cost, part of the cost is the process itself. The management of the keys not only the software that manages the keys but the property management of the keys to keep secure. And I think this online and offline sign may be it is not a good idea because I think it's just passing a bit of the best practice, I don't know. I'm not so sure because most of the recommendations in this kind of security is that you want something secured you must keep it offline. But it is not practical and it may be difficult to manage.

Ray:

Thank you and I'll keep this short. On the part of developing a security policy you're absolutely correct and that is what every Registry needs to invest time in. For instance, what do you do in

an emergency key rollover? What do you do when the HSM breaks down? All these procedures need to be developed.

Now also there is a solution for that and what Sweden is currently developing and actually Kira is the company that works for them as a consultancy company but they have developed this boiler plate policy that will make it very, very and we are using it and it will make it very, very handy to just fill out the variables that apply to your organization. So this will help you develop, for instance, procedures and it will also help to educate staff to the point that they're comfortable with the DNSSEC.

Chair:

Maybe I'll also make a small comment. The costs are only there if you do it the right way. Yeah I've got 2,035 names and I just take a postscript and do this and it works. Then Andre tells me once a month I forgot to sign .NA and so I must resign it and the cost of that was zero. But that is not, we haven't got a bank yet. Once a bank starts then we have to start looking at to do this right and according to the…once we actually use the DNSSEC for secure purposes and not just to show it can be done, then you start at a cost. But the point I'm trying to make even as a small CCTLD you can start getting the expertise and figuring this out and then step by step and when you have figured it out so far you can say okay it is now in a stage where we can allow a bank to use it, then you are probably easier in or more comfortable with spending money on it.

Louis:

Yes but I think the trust relies on the procedures or the process you follow when you do something. Then your clients or users can trust in you like you're doing right now. But if they want to check something or they have a little bit of not trust, then maybe the model can fall down. It is too easy to lose the trust. But keep the trust is hard because I think the expense and investment in trust is high in many of the things.

Sure we are not banks but we manage the URL of banks then we must be trustable.

Chair:

What I'm trying to say is it's easy and cheap to do it but it is not easy and cheap to do it right. As soon as we get somebody, we're doing it just because we wanted to do it. As soon as we get a demand and get banks to use it or secure companies want to have it, then we have to put all this in place. The point is start up costs can be small to get experience and expertise so then you don't have to make a huge investment up front. You get your learning curve, you figure out how this works and then when you start to put it into production where you have to put the expenditure you have that part of worrying about if it works or not, not so much.

Male:

My mind may have been wandering but I don't remember you discussing whether you're looking at N-Sec or N-Sec 3. I guess most people are looking at N-Sec 3 and if you're looking at N-Sec 3 are you looking at hashing everything in the zone or only those that you actually want to sign and then leaving the rest completely unsigned and unhashed? In which case your zone increase will be very, very minor.

Louis:

Talking about the bandwidth and size of the sign it is not my main concern. There is a cost there of course but not my main concern.

Male:

Bandwidth is cheap.

Andre:

I just wanted to comment the study of INISA because as far as I know it was concluded based on the queries of large CCTLD's that did GNSO recommendation and usually all those costs is the development of the system. So I think the 250,000 Euros it was written by us I guess and it is the lowest cost probably and it covers the development of the Fred system, the change in Fred system and all that stuff. So if you use some system like Fred or Coca which is done by a 3rd party you don't cover those costs of validation for small Registries are much lower. I don't think the small Registries should be afraid of the enormous cost of implementation of DNSSEC.

Louis:

How much could be the cost, the real cost? For example, use Fred.

Andre:

It's hard to calculate now holding the microphone but honestly the majority the sum went into the development of this system and writing papers that if you use really system like an Open DNSSEC, for example, I think it would be a very small portion of the cost. And remember also the cost of the (1:51:57 – inaudible) are very different then the (inaudible) cost in small countries that have small Registries. So really I believe it can be a small fraction of that sum if you're involved in a small Registry.

Male:

I was always worried about the hardware cost but if I can get it done for 1000 Euro it's not that much. If that is good enough for .UK it is probably good enough for .NA as far as I'm concerned.

Chair:

If there are no more questions I thank you very much for your presentation.

My son told me that in Germany he can hear us very well and so that seems to be working nicely. The first speaker after lunch is Roy Adams and he will roll over and die for us.

Roy:

Thanks. Earlier this year operators of the RIPE NCC noticed an increase in traffic to their main servers. What they thought was an apparent attack turns out to be something that has the potential to grow into a perfect query storm. If that does, we will have the equivalent of the China Syndrome. Now I'm not an alarmist, this is not theory, this is practice and it happens right now.

My name is Roy Adams and I work for Nominate and I did this work together with George Jeff Bettrick and by the way we just today deployed DNSSEC in the UK. So this is the graph I'm talking about. On the left hand side you see about 1200 megabits per second coming into the RIPE NCC name servers. In mid December it more than doubled. But what you see on the right hand side is a traffic load of about 20 to 30 megabits per second.

So when they looked at the actual traffic itself what they noticed was this is all DNS key queries. Now this a quite a clever attack because in an amplification attack you make use of the fact that responders are a lot larger than the queries and DNS keys are among the largest responses you can get from a name server.

In total, and this is just a simple calculation, you have about 980 to 1000 bytes for a DNS key response times the minimum we saw here, 2000 queries per second, it's a whopping 15.8 megabits per second. If you take this in order larger and you take a small CCTLD it will basically wipe it off the planet.

So who does this? What you see here is a graph of distribution combined with loads. On the left hand side you see a nice, long, tall blue line. Now this is what you expect, this is 10 minutes load of traffic and notice about 1200 unique IP addresses being involved in this and about 700 of them will give out about 1 to 2 DNS queries in a 10 minute timeframe. Now they run a lot of zones on the name servers, so 2 to 3 for 10 minutes that's not unexpected.

But on the right hand side you see the tall bar, the interesting part on the right hand side is this little blue line is about 60 unique IP addresses being responsible for 2100 DNS key requests in a 10 minute timeframe. This is per IP address so you need to multiply 2100 times 60 and you have the total amount being sent. So this is about a half a million DNS key queries they sent.

So what was so special? What was so special about mid December? Why did people try to attack RIPE at that moment? If you type in the 16th of December on their webpage, on the RIPE NCC webpage, you come to their DNSSEC deployment page and in this page there is a small line and I'll highlight it for you, it says, "On 16th December 2009, the current keys are removed…" what

that means is this is the end of a rollover period. Every June and December they remove a key and every March and September they add a key.

So in September a new key was added and in December an old key was removed. So this might actually not be an attack at all. This was a simple misconfiguration. Now why were so many clients misconfigured? The answer came about one month later, there is one operator in Australia and he noticed about 240,000 log messages in his log in the period of 24 hours. That is about 10,000 per hour, about 3 per second, I can't type that fast. This is enormous and he was also complaining about the enormous amount of bandwidth that was chewed up and you don't want to know how much CPU.

When he looked a little bit further there is this small DNS (5:18 – inaudible) that is part of the Fedora distribution and the DNS tool came preconfigured with the RIPE keys. So you have a bunch of static keys that you ship with the distribution and instantly on the 16th every one of these keys became still. So this is not that smart. This is about the same, at least in my mind, this is about the same as I remember this period where a large vendor of equipment have static IP addresses for NTP servers and this is just as bad.

So after the 17th of January, this was not that quite known, but RIPE sent an announcement that this was a problem and we hoped that everyone instantly fixed their resolver and, of course, they didn't but you see a small decline in load here. It goes from 2000 on February 3rd and this is just from last week and it's now currently just under 1000 queries per second. So this is still going on and even though it declines a bit it doesn't decline fast enough to mitigate the next role and I'll show you that here.

So here you can see this was not a one off event. In June they did the first roll and removed the first key and then in December and this is what I'm talking about, in December they removed the 2nd key and you can also see from this graph they added a new key as well but that was not so significant. You see only a small bump here. But you can see that it doesn't decline fast enough for the next roll.

So was this a one off event? Obviously, if it happened to RIPE it will happen to other deployments as well. And it did, it happened in June 2009 in Sweden. It happened in June 2008 in Sweden and you can see that here. And even though this is 60 to 80 queries per second, this is about 1 megabit per second and 1 megabit per second for…so even though 1000 queries per second or sorry about 60 to 80 queries per second is about 1 megabit to 1 ½ megabits per second.

The interesting part of this is that it was just 2 resolvers and the first one got fixed here and the other one got fixed there. So 2 resolvers being responsible for about 1 ½ megabits per second that's impressive. Now why so many queries? Resolvers are supposed to cache responses. I mean including DNS key responses and including if validation goes wrong it should cache that fact or even when the DNS key doesn't exist it should cache that fact. I mean resolvers are good in caching.

In general, resolvers should be nice because you have a lot more resolvers than you have servers. For instance, at Nominate we have about 11 to 13 depending if you count the IPV-6 addresses as well. You have an enormous amount of resolvers out there. So we did some research and it turns out there is a fundamental bug and it's what we call debt first search problem and it has to do with trust anchor validation, sorry the chain of trust validation and I will give you an example.

Normally if everything goes well you have a trust anchor configuration let's say for root and by the end of this year some people will have figured the trust anchor for root. So this is not theory. Then you get a record in; say an address record for www.DNSSEC.se and a signature and you want to validate that so you fetch the key. You want to trust the key so you fetch the DS and you want to trust the DS so you get the key and so on and so on. When stuff goes wrong and a trust anchor is off you see (9:51 – name) goes fetching every possible path it can get its hands on.

So in order to calculate this you have to multiply all these numbers and you get 600,000 queries, as fast as this can. Now the problem with (name), well it's fairly fast so it cannot do 600,000 queries per second. So basically what will happen here is your bandwidth gets eaten up, your memory gets swept out, your CPU is completely loaded and so you probably hit some other maximum first.

So we reported this bug early February and we got a very nice response. They acknowledged the problem and they told us the fix to this is actually a fundamental fix because it touches the core of the (10:38 – inaudible name) software. So it's a 9 version and so when they give out a patch they want to give out a patch to all possible servers. But then 3 days later they released (Name) 97 with a bug. Now this is interesting because (name) 97 is the first version that can validate the root.

Now remember the root is signed with a new algorithm called RSA SSA-2. And other versions of (11:11 – name) up to that point cannot validate the root but this version can and it has a bug. When we talked to them about this they basically said you need to exercise caution. I mean if you put an end to guns probably no one gets show anymore, right? Anyway it also ignores the lame DS problem because if you have a lame DS, a wrong DS, a faulty DS in a chain of trust you will trigger the same problem even though you as a resolver operator has done everything correct.

Then just last week it got a bit worse because now they released (Name) 962 and 962 is the most popular among validators. So people who are currently validating are using (name) 96 and (name) 962 now has the root validation back ported into it. So it can validate the root, it has this bug and absolutely no support for automatic trust anchor rollover. So in June/July if this thing is not patched people will configure a trust anchor. They might forget to roll it over by hand because it doesn't do it automatically, they are validating a root and stuff goes wrong and yet they have this enormous query storm.

So they announce a patch as soon as possible and I believe they are actually working on it very, very hard. I expect a patch within 2 weeks now and the people are currently deploying 970 and 962 and so I hope it's rather fast than not.

So let's talk about a perfect storm for a second. Currently we have DNSSEC deployment at the root. Guess what? It has a lame trust anchor called the Durs Key and it's an obscure key and they do that for a good reason. This is to avoid people configure the trust anchor and the resolvers. But here is what happens when they do, this is from a lab environment and this is only a few seconds of a query storm and this is a massive amount of queries it will generate. Please don't test this at home.

So going back to this automatic trust anchor role, automatic trust anchor role is something that is fundamental to the deployment of DNSSEC. Root is using it and this is basically a method to go from old trust anchor to a new trust anchor without any hands on. The 970 implementation is a little bit buggy and they promise a fix in 971. I mean it works but not when you're using views. But 5011 implementation, automatic trust anchor implementation is not planned for 962. So then you have DOV mishaps. This is something completely different. DOV is basically a DNS look aside validation but in essence it's just another chain of trust. We've seen mishaps in the past for this as well. First, when Puerto Rico rolled their trust anchor they notified the ITAR, ISC blindly copied the ITAR into their DOV but they do that once a week. Of course, when things are not correlated the updated trust anchors were not in the DOV Registry. And for about 3 days DOV users could not use PR.

We've tested this in a lab environment and this also caused a major, major query storm because it is just another trust anchor that's gone lame. Then we have another problem, this is the multiple source anchor problem. When you, for instance, have the Swedish trust anchor configured or the UK trust anchor configured you can do that now today, then even though you have a root anchor as well the TLD anchor will trump the root anchor. So this includes an invalid TLD anchor and it also trumps a root trust anchor.

So here is my doom scenario, right, let's say a few months from now the UK registers their trust anchor in the root and now we use the root to roll over the anchor. If people have configured the UK trust anchor and the root trust anchor and they forget to update or to remove the UK trust anchor, validation for UK will fail. I think this is also one of the fundamental problems, one of the phenomena that can cause a perfect storm.

So now we have this series of unfortunate events. We have things like DNSSEC.com and Fedora and you have the DNSSEC deployment at root, the multiple trust anchor problems, you have no automatic trust anchor rollover and optimistic DOV scraping. And to be fair, ISE doesn't do optimistic DOV scraping. ISE copies the trust anchors from the ITAR. But you have other DOV Registries that will simply just take your key whatever its worth and put them in their DOV deployment.

But there is something else. What triggers this problem is rolling over a key. A lot of people seem to suffer from frequent key rollover syndrome. They really, really want to rollover that key and some do it once a month, some do it at half a year, and some do it every year but it seems to me that people want to try to rollover their keys as often as they can. This was actually in 4641 in the ROC and is now being rewritten and what I was talking about before.

But the advice is completely misguided because it was based on a theory if you create a lot of signatures with a single key that it will eat the key. It is proven not true, that is not the case. You can generate an enormous amount of signatures with the same key without even coming close to leaking bits of the key and I'm talking about RSA here.

Then another intention to do this is to mitigate fallout from a broken key, a compromised key. Basically that story goes something like this if you roll the key every year and say after a half year it becomes compromised it only becomes compromised for only a half year because 6 months later you rollover again. But that is also not true because there is no perfect security. What that means is you use the old key to vouch for a new key. This is how 5011 works. You sign the new key with the old key. If the old key is broken, someone else can introduce a new key for you.

Again, that's not a reason to rollover a trust anchor. If it can be compromised in one year, you can do it in 6 months for twice the cost. Now it doesn't really scale that well if you take it to the extreme but for these numbers this is correct. Then if you talk to a lot of folks about this and I've tried it, they come up with different reasons. One of these reasons is we need to educate our operators. We need to test these procedures. You're not going to do this on a live system.

Can you imagine going into a nuclear reactor and saying let's change the rollover of the plutonium? Why not? If it fails we do it next month again and again and again. Well the fundamental rule we have at Nominate is core infrastructure. You don't just go touch that. If you want to test you do it in a virtualized environment. If you want to practice procedures you do that on a Beta system. And if you want to actually see if this thing works for real you test it on a ZS case, the zone signing keys not KSK's the trust anchors.

So all of these reasons to me are irrelevant. So what is the solution to this problem of the perfect storm? We'll fix the software already and stop releasing software that has known bugs. In my mind, ISE could have waited a little bit and I don't know why they stuck so rigorously to their release schedule. Another thing is to help ISE to basically spend more time on this and fix this. They need help with their (19:32 – inaudible) development. So if you're not stressed for cash, spend some cash on ISE to help fund (name) 10.

Then there is key rollover and don't roll keys as often as you can. Preferably don't roll at all. Take for example I think it's called root certificates in browsers, they don't roll that often. They are used for much more fundamental currently then DNSSEC is.

Then if, for instance, you're a CCTLD and you're about to deploy the DNSSEC you don't have to wait until the root is signed to make it official when the root is signed and put a trust anchor in the root. Don't use 5011 because it doesn't work quite as well yet. Don't use DOV, don't use the Orator, don't use web pages with listed keys because as I've showed you this can all go wrong. Just use the parent. If you're not able to use the parent because maybe it's not signed, only then do I recommend using 5011 or a DOV.

Now before I get a question later on let me tell you why Nominate actually rolls the keys every 3 years. We have a non cryptographic reason for that. We use something called an HSM and the term has been used today before. An HSM is a hardware security module and just like any professional organization we write off our equipment every so often. So basically we project that we can use this HSM for 3 years. Now the golden rule in HSM is you don't want to ever have the key of the HSM, ever.

So we basically roll to a new HSM and only therefore we roll the keys every 3 years. That's where the 3 years comes from. If we could get away with every 10 years we would do that.

So this is my presentation. You can read the full research on this web page but I also have another question. If you have deployed DNSSEC and you have done a key rollover and you have statistics from around that time I would love to see them. Thank you. If you have any questions I would like to have them.

Chair:

Thank you very much. That was an interesting presentation. A little bit over my head but I now know that I don't have to die anymore. Basically what you're saying is you only have to roll the key when you have to roll the key.

Any other questions?

Warren:

So just because you depreciated the cost of the HSM I don't understand why you then need to throw the HSM away? Like you don't throw your servers away every 3 years.

Roy:

We don't actually throw the HSM away but the idea is we can't predict the future and we don't know if that same HSM is supported 3 years from now. So we stuck to the commitment we can get from the HSM vendor and they don't go further than 3 years from now. So if we can't get support we don't rely on that.

Obviously, we can buy some more HSM's and stick them in a closet for the future but maybe by that time there are better HSM's, cheaper HSM's and faster HSM's and maybe those HSM's can't import the key from the old HSM, so it's just safer and easier to roll the key.

Chair:

The next thing on the agenda, Andre Phillip will speak a bit about experience they have about DNS attacks, research they have done about it. It is not necessarily like I thought initially that he would present an attack and how to mitigate it but he will talk a little bit about some research they did.

Andre:

Thank you very much. My name is Andre Phillip and I'm from Dutch Registry from Czech Republic. I would like to talk a little bit about one study we are doing in our Department which is called (25:31 – inaudible). It is all about DNS cache poisoning and how real this attack is and what are the conditions to perform the attack. We did some testing in some real scenarios.

My agenda will be a little bit about theory. I know you are all great experts in that but I still have some slides about DNS and I apologize for that. I will talk a little bit about the cache poisoning but I just have 20 minutes so don't expect some very deep dive into the theory. I will just mention some highlights of that.

I guess everyone knows this picture. I took it from Wikipedia. The funny thing is there is a bug in the picture on Wikipedia pages. But you all know now DNS works and that it has to go through different stages to get the right response.

This one is more interesting. It is about the structure of how DNS works. So between the operation system and the whole DNS system is something called DNS Resolver and it as a local cache. The whole attack is about poisoning and inserting bad data into that cache of the DNS Resolver. There are many ways how it can be done. So if you ask your resolver from some information on a DNS query it resends this to the other DNS servers. And for an attack code to be able to poison a cache one way is to send exact replay expected from the authorized name servers.

To be able to do this it has to know a few items in the query and it needs to be repeated in the response. Those items are the source address, which is usually known because it's the address of the DNS resolver, the source port which should be random in many implementations around like in this room and it's pure random. It's about 16 bits of information. Then the destination address and it's usually known, it can be more than one. It can be 2 or 3 authorized name servers but usually it's just a small number of choices. And the destination port, which is well known and it's number 53.

One important thing is the query ID which is also random ID and should be at 16 bit. And, of course, the query section but that can be somehow sent by the attacker so the attacker knows the query section. So basically if you look at the table what you need to guess is about 32 bits of

information maybe a little bit more if you count also the name of authorized name server destination addresses. So it's roughly 33 bits of information, which is quite a huge number.

Another problem is the fake response has to be delivered before the regular one comes. The attacker has to be faster than the regular response. So how you can do that? Under picture you'll see on the screen there is a timeline on how the communication can look like. Here is one way how you can poison the cache.

So you first send a DNS query to the resolver, you as the attacker and the resolver starts working for you. So it sends a DNS query to some authorized DNS server. Then the attacker starts to send fake responses to that query, so he tries to poison the cache. Those are the blue and red arrows. As you can see, some of the arrows came a little bit earlier then the DNS resolver expected and some of them came a little bit late, so those are blue. Those that really hit the point but we don't know if they were successful or not because we don't know if the attacker was able to guess the 32 or 33 bits of information.

They hit the DNS resolver and they are red and its so-called attack window we call it that way. So that's the time that the DNS resolver can be attacked. All queries that came after the regular response are blue again. Then the resolver sends response back and then the attacker knows if his attack was or wasn't successful and he can repeat it.

So there is some window which is the red area which shows the time you are unable to attack the resolver. So you need to wait a little bit. You can effectively use the time for attacking to the resolver and then you repeat the attack again. And there is a certain probability that you will be able to forge the first response and you will poison the cache.

Once you do it you can set a huge GTL so you can hold that resolver for a certain time so that it responds the way you want it to respond.

So that was the theory but on the other hand before Kominsky came with his idea, everybody expected the window for attacking the resolver is very short. When it gets a valid response it keeps it for a certain time in the cache and the window for attack opens after the GTL expires. But Kominsky came with the idea that this is not actually true because you can create some sub domains of the domains you want to attack and you can use sub domains that really don't exist that the resolver has to ask the authorized name server for the response. And you can forge the response in the authorized records and additional records you can forge the false information about the domain.

So he really found a technique on how you can repeat this attack quickly. This became an issue for resolvers that didn't randomize the source port of the query. It was just 16 bit and you really need to guess as an attacker. So it was easy with Kominsky knowledge to speed up this attack.

So the direction was that everybody started to use portal randomization, ID randomization with very good random number generators and now the question was and is how is this secure? How quickly can you attack such name server?

What we tried was to start a really brutal attack and send all those fake responses. We generated all possibilities of source port and query ID from the possible ranges which is all those 32 bits. We used some modified implementation of (33:54 – inaudible) and also we did some measurements with flooding the authorized name servers. So we wanted to try and find some ways on how to improve the attack and see the limit of how quickly we can attack the resolver.

Now I will try to go quickly through the theory. I know that if I have any firm line in the presentation that lowers the attention of people so I'm sorry for that. I know I will lose $1/16^{th}$ of the people but it is necessary to understand how it works.

So the time of attack that is very easy is just the number of windows you need to have for the attack and the width of the attack window which is in milliseconds and it's the width of the attack window but also the overhead window. So you need to count both actually. If it was letter N in the (34:58 – inaudible) which can be written like that so that shows what number of attack windows we need to be successful on certain probability levels. And you need to set up a probability level on your own and so usually use values like 95% or 99% is a probability attack so nobody can be sure that it can ever be performed. But there is a certain probability that it will be successful.

NP is the probability of guessing ID port and destination address, which can be described like that. It again is very easy and it's just a number of queries in a window and this can be measured. Then it's multiplied by the old let's say the bits like number of ID ports and number of authorized servers you need to guess in the attack.

But the most important thing is we know D U & S, we set queue as we want and let's say we want to work on the probability level of 95% so that's an S and there are 2 values we need to measure to know how this attack can be successful and it's F & W. so it's number of queries per attacking window and the width or length of the attacking window. So how much millisecond this attack window is open plus the overhead.

So again we just measure 2 very simple units and don't be afraid of all the complicated math. And that brings me to the testing scenarios. First of all, it is called testing scenarios but really it was made in a real network. We just really bought a regular hosting and set up a regular resolver and so it was across the internet exchange point on real networks by real internet service providers. So it is testing scenarios, however, it was really performed like it would be in the real network.

We used 2 authorized servers for the attack. We tried to forge queries and there are some numbers like what was the average DNS message size and the ports and ID's we tried to guess,

so almost all of them. It looked like that, for example and this is scenario 1. The attacker has bandwidth about 100 megabits per second. The server was connected to a bit higher speeds; let's say 1 gigabit, definitely more than 100 megabits. And the type from the server to authorized name server 10.843 milliseconds. So it is on very fast networks the attack was quite hard. The equipment of the attacker is quite obvious and not very expensive. And as I said the window was quite short and as you can see there the response to authorized name server is very shot.

What we saw, we measured that the width or length of the window is about 1 millisecond. We were able to forge about 60 queries per that window. And to be able to do this we generate a stream of fake responses about 60 megabits. So that is one thing we need to keep in mind, the stream was quite visible. The interesting thing is the window is just one but the overhead time is 10 milliseconds and that means that you lose about 90% of your time just waiting for opening the window.

Interestingly, the theory says that if you work on the level of 99% probability you are able to attack this resolver in less than 90 days. So even in this quick scenario which seems to be quite impossible to attack you are able to do it with certain probabilities. So let's say after 60 days you are really successful so that was one scenario.

Another one which is from a more European point of view and very typical when you have name servers in the US and the queries need to travel below the ocean, you have like 1 70 milliseconds time to do authorized name servers, so this is much easier to, the window should be much, much bigger in this case. Again, the attacker as 100 megabits per second network so the same condition as before.

And as you can see, we were able to get something like 100 more queries per window, so the speed of the attack was something like 100 times faster. The stream of the responses were roughly the same. The window width is very high in this scenario and it's something like 160 milliseconds and the overhead is much smaller than that so it's not important in this case.

As you can see, to attack such a domain which is a little bit far away from you, you are able to do with 99% probability in less than 9 days. So on average you can do it in something like 3 or 4 days. So it can be done very quickly. But again the stream is quite high.

So we were thinking about how to improve the first scenario or similar ones to the first one, how to make the attack more efficient. We tried to just attack the authorized name servers. So they were flooded so they couldn't respond so quickly as they wished. So that makes the attacking window a little bit bigger and the attacker has a better chance to attack such domains. And we made a complete new set of tests, so in the 1st case we were able to send about 30 to 40 fake queries per second and the attacking window was very small, overhead was much bigger. When we add this attack the window width rose from something like 0.5 milliseconds to 741 milliseconds. It was a quite huge deviation and it varied a little bit but on the other hand the

window size rose dramatically. And, of course, the same applied for the number of queries we were able to send.

So in the first scenario we were able to send about 40 but with this attack we were able to send something like 50,000, so a thousand more and overhead was quite not important in the second case because it's much, much smaller than the window itself. And that means in this scenario we are able to make the attack shorter from 10 or 11 days to 3 days on the level of probability by 90%. So you could be able to make this attack in less than 3 days.

So as you saw this was very well connected, authorized name server, very close to your resolver. It should be hard to attack it but with this improvement you are able to make it in 3 days.

So that was the theory and we wanted to see whether this works in reality. So we set up a real resolver and DNS server and we wanted to attack. We chose domainexample.net in our case and don't be afraid the resolver wasn't used for clients; it wasn't an attack on some innocent clients.

There was one important thing from all points of presence, the rate time of one of the authorized name servers was much smaller, it was for BIANAserver.net so we expected the server, the use would be for the B-IANA and it was true. We measured it and it preferred the B-IANA in about 98% of the queries. So that lowers one bit of the information because you don't have to forge 2 responses for 2 name servers, you can use just one of them.

And we also ran this test for DNS resolver with no port randomization, the old one. Of course, as you can see from the table and you can choose the stream you wanted to send, if you want to send a really invisible stream of something like 50 kilobits you can do it in something like an average of 200 or 300 seconds, so this attack can be done very, very quickly. But that is in case you don't do the port randomization and I hope that almost everybody does today.

So we tried to attack the reclusive name server with randomization and we ran about 6 tests. As you can see, we tried different streams and different conditions. Again, as you can see from the times we were able to attack it and the time that was expected on the level of probability, so in one case we were very lucky and made it in the 15% of the time that was expected for the attack. But basically we proved we were able to make such an attack in the speed we calculated.

Again, this is just a presentation from a study and the study covers much more examples and much more theory behind it. So if you think I'm too quick I'm sorry but my time slot is limited. There are many more numbers behind those things.

Again, it shows it can be done in a real environment and it can be done in the time you can wave like attacking name server in one day, 24 hours, is quite nice. Probably you could do it before the administrator would be able to react. The resolver sent us wrong data after those attacks for the domainexample.net.

So what are the costs of such attacks? As I said, we did it in our country in Prague and we used 2 servers for that and we paid for 2 server hosting by commercial companies. So it was something roughly like $300 US dollars per month. We spent about 3 weeks on that but we spent 3 weeks on working on all scenarios and I think it was 6 or 7 scenarios in the document. We set up the network, we wrote the document I'm talking about and so in the real world the attack would probably be much quicker. Of course, you would need just one server and one server hosting, so I think you wouldn't spend more than $2,100 US dollars per attack. So if the attacker believes the domain is important the money is probably not the problem.

Just a few remarks on what can affect the effectiveness on the attack. First of all, the balance of authorized servers I spoke about A-IANA servers and B-IANA servers, if one is much closer then it's expected the resolver will use it more. So it's better to use this one to forge the fake responses. Then if there is a higher number of authorized servers, of course, it decreases the effectiveness of the attacks. Also, if you are closer to the resolver and have higher capacity than of course the attacker has a little bit more complicated job. That is a quite basic rule that applies to many other attacks.

Port and ID randomization you can test your resolver. An example on how to test your resolver whether it uses the real randomization of ports and IDs and again bandwidth of the attacker. And last but not least is the monitoring because if you are really looking at your resolver you could probably see some change in the data flow, so you can guess that it's under attack. And, of course, the DNSSEC.

So I am close to the end of the presentation, so after Kominsky discovered the problem and many DNS's were patched we're still not over the problem. It still exists and you cannot do it so quick like before Kominsky, like in seconds but you can do it in hours. So DNS you can try a lot of techniques on how to make an attacker's life harder but he can still be successful. You cannot avoid it completely and the equipment for such attacks is very cheap and it can be done in hours or at least days.

So that is all about the study. As we are one of the supporters of DNSSEC, we have more than one year of DNSSEC signed domain and we have about 80,000 signed domains in our zone which I think is the record today. My advice is try to implement DNSSEC if that is feasible for you.

I have one remark on my final slide; the study will soon appear on the pages there. It is not in Czech language only and is probably not very helpful for you and so we will translate it and put it to those web pages. So thank you very much.

Chair:

Thank you very much. I'm not sure I understand all of it. As soon as it's translated into English please let me know. Any questions?

Male:

Andre, what sort of campaign are you guys running, if any, to encourage uptake of DNSSEC amongst your domain holders?

Andre:

The campaign? Well there are many ways but honestly what's worked best for us and I think Pablo which some are trying to hide behind has a presentation about it on Wednesday so we will see it. But really we are trying to communicate with the Registrars that makes the majority of the work but, of course, we already tried to push the topic to the media. We were in the main news and TV stations in the country so we were quite successful to be able to sell the topic in the media.

Male:

I would like to add that a part of the campaign was to show this kind of attack on the conferences and show to the audience it takes 10 minutes to run it and then you just wait for the jackpot.

Chair:

Any other questions? All right thank you very much. So this pinging in the background seems to indicate that our remote participants are dialing in. can they hear us? Can they say something?

David Dagan:

Yes good morning.

Chair:

Who is this?

David:

This is David Dagan in Atlanta. We also have Chris Davis in Canada on the line.

Chris:

Hello.

Chair:

Good morning to you and good afternoon from here. Can you see this on the Adobe?

David:

We both have some visibility on the slides and I'm looking at the tail end of a presentation DNS 2010-03-07.

Chair:

I've put the next one up, can you see that?

Chris:

Not yet.

Chair:

Anyway this is Chris Davis and the presentation is online and you can just ask me to advance the slide. On our screen we're seeing the Defense Mariposa Briefing. So you can start at any time.

Chris:

Okay just would like to wait for the slide to finish so I'm not confused to where I'm at. All right is everybody in the room and we're ready to go then?

Chair:

The screen on the wall shows your presentation and we can see it now on the Adobe as well.

Chris:

Perfect. I want to get through this fairly quickly because David Dagan has a slide deck after mine that fits in together and we didn't have time to combine the two. If we can move to the next slide and I can do a quick introduction to Mariposa.

For those of you listening that haven't heard of the Bot.net we discovered it back in May of 2009 originally. I'll go through at timeline fairly quickly. I apologize if you miss something but these slides will be available to you and I'm always available to answer questions later.

The current state of Mariposa are 3 arrests have been made in Spain. We can't really talk about whether more are going to happen but it is probably safe to presume there are more people involved and people are looking for them. So far recovered off the bad guys hard drive, well it's actually 5 or 6 different hard drives that still need forensics performed on them but on just the first one that they've finished performing forensics on, there has been over a million personal credentials recovered thus far.

Since March 3rd there has been 903 international news stories on this bot.net. The news really seemed to like this one. So it's getting a lot of traction and that could be good for us to get policies changed in certain countries and it helps raise awareness and that's one of the biggest

things we have to struggle with as security people. A lot of times people don't really understand what a bot.net is and raising awareness is really good.

The last point there, this is one of the things that always blows me away. From December 23rd until last night at about 10 p.m. my time here in Ottawa there have been 15, 550,000 unique IP addresses connecting to the Mariposa sink hole. Now what that means is when a computer that is compromised with Mariposa beacons are called home, there are 2 variances that we call Mariposa A and Mariposa B. Variant 1 has a 7 byte payload, Variant B has a 22 byte payload and we actually look for that exact payload. So if somebody port scans our sink hole, we're not going to count them as somebody who is compromised with Mariposa because they have to send that exact payload, which is the kind of hello I'm awake payload from Mariposa. So sending those payloads has been over 15.5 million unique IP's.

Now one of the questions we get asked all the time is what does that mean in number of computers? I don't know the answer to that. One IP address could be a nat point with 100 computers behind it obviously and 100 IP's could be one computer that is connected to a dial up or DSL line. So the numbers are almost impossible for number of computers. You will see in the press some people claiming its 12.X million computers. That wasn't us stating that that's other people sort of skewing numbers around.

This slide is the beginning of the Mariposa FAQ. What does the mailer do? Well the mailer does whatever it wants. When this slide was first put together it was for a slightly less technical audience than I'm talking to right now so I apologize if you're bored already.

Anyway the idea here is that what everybody has thought of as viruses they don't really exist anymore. The I Love You virus, the Melissa virus, things like that. It actually hasn't been viruses like that written in a couple of years now. It is all financially motivated malware and it's designed to get on your computer, it's designed to hide there, it's designed to hand over full control of your computer to the bad guy and AV is failing to detect it and we'll go into that a little bit later.

This particular piece of malware was built on a Butterfly Black Jet, which is primarily an information stealer. The first thing it does when it connects to your computer is grab all your stored credentials in your web browser and pass them back to the bad guy. The second thing is does is it drops a second stage key logger which also and I can't remember the exact time, I'm going to say daily basis would dump the key log data off to a secondary, another site for the Bot master. So those were the first 2 things it would do. But it certainly could receive any command given to it.

So who is responsible? Currently there have been 2 or 3 arrests made in Spain in relation to Mariposa. The primary Bot master, I guess as we want to call him is a guy named Nakario I think was the handle he used. He was a 31 year old sort of an average guy apparently from what I heard from the Spanish police. Not highly technical, bought a kit and that's all he needed to do

and it's a pretty graphical user interface to build the bot and there is not much too it. Kind of scary really.

How does Mariposa spread? Well by default the malware is designed to spread through instant messenger which is actually MSN messenger. It will send out links to everybody who is in your MSN messenger list. The USB key thing although not new and not interesting, the thing that is interesting about this within Mariposa is we've been monitoring certain CNC domains from Mariposa further back, before December 23rd when we turned off all the main command and control domains. Sort of going back further we had a few of the domains we've been looking at and what we were seeing is a pattern of reinfection. We really think this related to the USB key thing.

So if you have a really old Mariposa variant to get picked up by your AV let's say and get wiped from your machine, but maybe it was on a USB key from a couple weeks earlier when you shared a file with a buddy, then another computer inside your network or your own computer gets reinfected again depending on the way your AV is set up. So we really did see from the same networks and the same companies patterns of reinfection where it would go away for a week or two and then come back, so fairly effective.

Okay why did we call it Mariposa? Well it's based on a butterfly kit, the bad guy was in Spain and we figured that out pretty early so we called it Mariposa.

What companies and organizations are compromised? Obviously we want to be very careful talking about particularly who it is but the sink hole does give us the granularity to determine what networks have compromised systems and we're working with everybody and authority we can to try to let as many companies and governments know. It's hard because there isn't really any central place to give this data to and I have a small company in Ottawa and I can't have everybody in my company calling up different government departments around the world every day. So we don't really have a good solution for this and it would be something that I would be really interested if some of you guys have ideas. I would be really interested to hear that.

This is just a map discussing the size and we already talked about like I said it's unique IP addresses and it's just a little world on fire map. Dave Dagan can speak to that later. Moving to the next slide please.

So this is where this conversation really comes in, how do we detect Mariposa? Well we detect it using GNS. We're looking at a friend's sort of authority data who owns some domains and we noticed there were some really interesting traffic patterns in there. And as part of what we do with our behavioral engine analysis we developed at Defense Intelligence and we started to look at it and say wow this is not a bit torrent site, this is a bot.net. That's how we found it. We never even saw the binary until a few weeks after we started tracking it.

So DNS is really effective against malware. I talked about that on my last slide but it is really essential. So having, right now we have a few sensors with a few friends and research partners where we see some DNS data from different parts of the planet but sharing DNS data even when its anonimized which is fine is really important to find new stuff.

I'm going to skip now and go over to the timeline. So can we move to the next slide? So going back to January 2009 when the butterfly kit first hit the web, Rafe at Spy which is now HP, he's an incredible security researcher, very, very good application guy. He's not into malware, doesn't really know much about it but he saw this come out and was like oh that is not good. Like I said, in May we discovered it and started tracking it and in June/July we noticed it was getting really, really big.

If we move to the next slide, so in August/September we realized this thing was really starting to his some critical systems. The thing that was interesting about it at this point by only looking at DNS data of course, it appeared that this was really targeted after government and financial and Fortune 1000 corporations. It seemed to be giving home users a pass, which we found very curious. We think now that we actually have some granular visibility, we find that is not the case but at the time it looked that way. So we were wondering maybe if he was splitting home users off into a sub section of the bot.net that we didn't have visibility into on the GNS side, which we did get later.

So we started to notify organizations that we noticed to be compromised. If we move to the next slide this is responses we got from those people. Now this is responses we got from banks, from government agencies, people who should know better. This isn't me calling my grandma and trying to explain something to her. It was so frustrating and again I think its user awareness and I think it's good that Mariposa is getting all this press, although I have not time to do my job.

So the reality of it is when we were contacting people this was status 41-AV vendors were detecting the binaries we were looking at, at that point. There is now close to and I can't remember exactly how many but let's say 800 or better variants or unique ME5's associated with Mariposa. But this time very low detection rate even though the banks we talked to said absolutely not and we're going to sue you if you talk about it at the Canadian Banker's Association which was the governing group. They came back and said oh yeah it is in the banks and we're taking care of it. The World Bank released a snort signature for Mariposa, which was really nice of them and Palo Alto came in with a plug in for decrypting the traffic between the compromised system and the CNC.

So in October/December we started taking over connectic control domains and in retaliation the bad guy registered Def Intel which is our domain, DefIntel.com and he registered DefIntelSucks.com, .net, .org. so we really started to track the bad guys at this point. The people that were part of the Mariposa working group that we had put together were really helping to sort

of dig in and pull out email addresses and things like that. We were able to cross correlate that data. Of course all of this was to gather data for law enforcement.

Then on December 23rd our small working group executed a time to take down all the bot.net CNC domains. We had written this fake obituary for the newspaper and we were going to put it up on our blog but we were told by law enforcement don't talk about it yet, so it never went up. Actually I think it's up now but I thought it was funny so I included it in the slide.

Now on January 22nd an entry level employee for a European Registrar was bribed by Nickario for 500 Euros to get one of his about 23 connectic control domains back. The bribe worked and the employee gave him the domain back. It took us 2 days to get the domain back away from Nikario, the bad guy, and so during that time he was able to rebuild part of the bot.net and update them with some new binaries and add a couple of new domains to the pool of what needed to be connectic controlled. So we were trying to track that and pull the domain back away from him. We did get the domain back and the employee was fired and within 5 minutes of getting that domain back, here in Ottawa we were hit with 900 megabit UDP flood against our sink hole infrastructure. Which is fine and it's our sink hole infrastructure and we expected to get attacked and it doesn't really affect us in our day to day business.

The problem is our fiber provider only had a gigabit link going into Ottawa. So it took out all the fiber providers customers in Ottawa, which included some universities and some government agencies for several hours. So that was kind of less fun.

So on the 26th of January the (1:10:59 – Spanish name) and FBI issues a search warrant. They arrested the main guy. They seized his computers and his hard drives. Apparently Guardius Seville has publicly said apparently the bad guy was quite scared and sat down in the room and said oh I'm going to tell you everything but I want to talk to my lawyer first. He talked to his lawyer and then came back in and said hey I found that computer on the Metro, it wasn't me. He thinks that is his defense and unfortunately it may be. It is so hard to prosecute these guys sometimes.

On March 3rd Spanish law enforcement holds a press conference and we released the updated technical analysis on white paper, which is available on our site if you go to DefIntel.com if you want to read the updated white paper.

Now we're going to scroll through a couple here and I want to skip the why didn't antivirus detect this. Again, this is for a less technical audience than you guys so if we can skip. Okay I just want to hold on this one for a minute. Even though I know you guys are technical and understand this but the interesting thing here is this virus total slide, this is really the current state of detection worldwide. Out of what 69,000 binaries 1,400 were detected. Now that is amazing to me. It depends on who you talk to and which stat you want to listen to but there are somewhere between 30,000 and 50,000 unique malicious binaries based on MD5 pushed into the

wild every day. There is no way an AV company even with 100 guys chained to desks writing signatures are going to be able to keep up with that.

So a real dramatic shift has to occur within the security industry. We are really losing this battle. I am such a proponent of DNS is the way to do this. Let's move to the last side please.

This is just my little soapbox bullet points. So I really believe that fighting malware is everyone's responsibility. I've heard from some DNS operators who say well we don't or domain Registrars, we don't police content or it's not our business. Well I really think it is. I think its everyone's business and I think unless you want the bad guys to win that you've got to take a stand even if it's just a moral one and say okay that guy over there beating up that grandma and stealing her purse, I'm going to step in or at least describe him to the police or do something as opposed to well it's not my job to police people. Yeah it is.

As long as malware uses DNS, DNS is one of the most effective ways to fight malware. So we need to use it. the Department of Justice in the US said that and I don't know if this quote is true or not, but no it was the Department of the Treasury in the US said that cyber criminals are currently making more money than drug traffickers or moving more money or something to that effect.

And one of the other things we really run into when we're talking to CCTLD operators and GTLD operators and some Registrars is that even if they will respond to abuse and shut down the domain, they would not be willing to redirect it and that hurts the ability for us to (a) find the bad guys like we did with Mariposa and that was absolutely all about being able to get that domain to a sink hole, being able to see who was compromised, where they were coming from, how it was spreading. All of that stuff is really important and we can't do that if we can't see the traffic that is hitting the domain.

And as the last bullet point there says is that sink holes allow us to truly understand a bot.net and you can't fight something you don't understand. And looking at DNS data is maybe the second best thing, so if a vetted researcher comes to you and says hey you know me through this person and I'm trusted, can you please shut down this domain? And if you don't want to redirect it to me can you give me a snapshot of what the DNS query volume looks like? Geographically where the queries come from? All of that is going to help up stop these guys and catch these bad guys. If we don't do it, if we just shut the domain down all the bad guy is going to do is go register another one.

Thank you very much and Dave Dagan is going to step in now and finish off the presentation.

David:

Thank you and I wonder if we couldn't get Christina to upload the 2nd set of slides. They have some interesting graphics worth looking at.

So thanks Chris and my name is Dave Dagan and I'm a post doc at the Georgia Institute of Technology in the Information Security Center. I participated with the Mariposa group in tracking and finding the Bot masters in the Mariposa bot.net. I wanted to present a couple of very short slides dealing with some lessons we learned in the process.

Slide 2 of 8 I summarized a few of the salient points. The Mariposa bot.net is characterized by having a very large number of victims, somewhere in the order of millions. It is difficult to measure an instantaneous amount but certainly an extremely large population set. We also found that international coordination was very possible and very easy to do. I thought originally when we started this project that it would be the hardest part but as it turns out many of the essential relationships that we came to rely on were facilitated by previous meetings that ICANN had facilitated. So this was absolutely essential for us to get started and it turned out to be something that was the easiest part.

The domain take downs were the hardest part. When we finally identified individual domains that we thought were malicious we faced 2 problems. One was resisting the temptation to just take them down immediately and finally when we had sort of an all clear signal from law enforcement that they were done with their work, the actual mechanical take down of these domains proved to be somewhat difficult in some cases.

To the first point, domain take downs in general are a cottage industry in the information security field. It is quite common now for information security researchers to identify malicious domains and then have them removed. There is often a lot of white papers and press associated with this. People have claimed they've taken down bot.nets.

In truth I think we have to recognize that this practice although it is useful in the short term, merely in the long term just simply inconveniences Bot masters. It only costs a few tenths of dollars for Bot masters to become active again if all you do is take away a few domains from them.

We approached Mariposa from a different way of thinking. We thought that if this is truly a cyber crime we have to track down the criminals and not merely take away a few of their assets. As Chris also noted, this is an ongoing case and in deference to the ongoing Spanish judicial process we will present just some statistical highlights in the remaining slides.

On slide 3 of 8 we see a CDF graph of the Mariposa victims by country. If you take, for example, the top 50 countries where victims are located that accounts for approximately 90% of all the victims in the bot.net. Now this is extraordinary. There are about 210 different ISO 3166 country code designations represented in Mariposa but the top 50 are really needed to make up about 90% of the population base. Usually with a bot.net it's about a dozen or so countries that make up 90% of the population base and usually Europe and North America.

But in this case it's quite diverse and it's spread out. Let's look a little bit further down at what that means. On the next slide, first to summarize the meaning of that CDF well it did have a very, very broad population base and we find the US is certainly a prominently represented victim population but the bot.net itself was not focused on merely the traditional North America and Europe infection base. We believe this may be due to 2 factors at least. The message base propagation mechanism that Mariposa was using and also the fact that the Bot masters were renting out the bot.net from time to time and this may have contributed to the victims over time.

I think there is some public press discussions to the effect that the Bot masters were earning about 3,000 Euros per month just renting out the bot.net on top of all the other identity theft and banking credential theft they were conducting.

The implication of this CDF is we face a common threat. This is not a situation where we have a particular internet threat originating from one localized area and threatening as often appears to be the case in the press a particular country or small user group. It truly is international.

The next slide on 5 of 8 we see a time series plot. What I've done here is taken a couple of weeks worth of sink hole data and taken the top I believe its 15 or so countries that were represented and plotted the number of victims in every 15 minute epic. I apologize and I actually intentionally made the graph somewhat difficult to read but you can actually see the number one property this is intended to convey, which is there is a very, very strong diurnal pattern to the victims. As you can see, there is a very large rise and fall and some of the victim population groups never fall below 50,000 and some go as low as 20,000 victims per 15 minute epic.

But as you can see there really is a shift between night and day. The implication of this is noted on the next slide, 6 of 8. And aside from the brief censure outage in the center of the slide you'll see that there is a very strong diurnal pattern and this means these victims are end users and not servers. Additionally, if you look at the key you'll see that the ISO 3166 CC designations are not the usual victims. We see a rather diverse group of countries where victims are found and many of these are not traditionally represented in other bot.nets.

So on the next slide, on 7 of 8; I've plotted various radii circles that indicate the number of victims in a given geographic area. If you recall in Chris Davis's slide deck he had a particular slide that showed a single red to where every single IP address could be found through geo-location. It indicated a pretty good distribution of victims around the world. But when you adjust the size of the plotting to indicate the number of victims you can see that this is actually not a traditional bot.net that has large numbers of infections in Europe or there are certainly quite a few in North America but is very well represented in Mexico, in China, India and so it truly has a very large international following.

On slide 8, I draw some inferences from all of this. Really this Mariposa is the inverse of some popular notions of bot.nets. The common belief amongst bot.net researchers is often that bot.nets are created in the East and the victims are always in the West. Well Mariposa proved that it can

actually be the opposite. In this particular case the Bot masters were in Spain and the victims were widely distributed, many of them in the East as well.

And the implication we take from this is truly this is a common threat and not confined in origins to a given geographic region. From our efforts, we've learned the following: international coordination is very, very feasible and we thank ICANN for its previous role in facilitating early relationships that proved essential for us getting things done. We found that domain suspension was perhaps the most difficult part. I know many Registrars have a standard where they require a court order or some very strong indication that there is some criminal mischief associated with a domain.

I completely understand the reasons for that and certainly it makes good sense to contain liability and if one were to open up a Registry to arbitrary suspension that vector itself would be used for malicious purposes. But I think we found in our situation an unusual scenario where we didn't necessarily have a court order but we certainly had a court that was prosecuting and still to this day is prosecuting a criminal case against individuals for use of selected group of domains.

While those individuals are in the judicial process, we would like the domains to be removed and remediated. We think that perhaps there is some other standard that can be developed and I myself would be very willing to go out and obtain insurance so that I could indemnify any errors or mistakes that were made and perhaps there is a method which bonded researchers with Errors and Omissions policies or proper indemnification could provide an argument for suspension of a domain in lieu of a court order.

I simply find that the court mandated remediation effort as a rule although it certainly makes sense from a risk mitigation point of view, I don't know that it is the most efficient way to help clean up the internet. I think we can still accomplish both goals of giving strong assurances to the operators that there are very good reasons for suspending a domain and limitations on risk and also allow researchers to make swift progress in cleaning up bot.nets.

One final point, part of this case at one point turned on the presence of a single DNS packet. We were tracking where this Bot master might be and we ended up using DNS at one point to triangulate the location of the Bot masters in Spain. We had other indications but this was conclusive proof for us. This also proved to be essential for engaging local law enforcement. I think the lesson we draw from that is that DNS monitoring is a very useful activity and it goes simply beyond epidemiological data but can also be used in an operational sense.

I do note on the ICANN's public comment website there is a proposal for a DNS CERT. I think there are many different types of mechanisms that could be used to accomplish large scale DNS monitoring. Since ICANN is currently considering the wisdom of a DNS CERT creation I think that is a very useful vehicle for doing just that. I believe the comment period is still open and won't close and off the top of my head I believe it's March 19th or so. I would urge people who are listening to this presentation to give a serious look at the proposal that Yuri Ito has put

forward and to comment appropriately on it. I, for myself, I'm going to be endorsing it and I think it's a great idea. I think our work in the Mariposa case is an example of how creative use of DNS monitoring can be used to help clean up infections on the internet.

I am a researcher at a university and ideally this would be an activity that becomes the daily chore of people working at a CERT. I would urge people to consider that proposal closely. I think it deserves a close read.

I think we may have some time left and I believe Chris and I are ready for any questions.

Chair:

Thank you very much and those were 2 very interesting presentations. As a matter of fact, Yuri has arrived from a previous presentation so she will sit on a panel and give us a short presentation. So if you want to remain on the line and listen you're more than welcome.

Are there any questions from the audience? As I said this morning, they're all in awe. All right thank you very much to be available on such short notice. I will correspond with you in email because it was also short notice that we haven't really introduced ourselves. So we should stay in touch. Thank you very much.

Okay so the next or last item on the agenda if I'm not mistaken will be our Incident Response Panel. Steven (last name) will moderate it and we had invited Yuri Ito who will give us first a presentation on DNS CERT, Patricio Pablita has sent us his presentation and so if he's dialed up he will give us a short presentation about what happened in Chile in particular with a view from his own operations. Then the Chair after the Incident Response Group is available. Then we have Roy Adams who is sitting on the panel.

Steven:

When Dr. Eberheart asked me to moderate this panel I asked myself what exactly we were trying to cover in this session. And until September the term Incidence Response to me invoked a notion of having to respond to one more cyber threat against the Registry, to deny all attacks against our servers, continuing attempts to harvest registration information via assaults on the WHOIS server and the usual mischief one sees against registration websites for those of you who run those as well.

The appearance of the Conficor Worm and the subsequent appeal from ICANN to managers of CCTLD's that were identified as targets of that worm added an additional dimension to the notion of incident response. ICANN's ability to respond to this threat was both ad hoc and haphazard. We rely to a great degree on outdated contact information, implicit trust and the ability of CCTLD managers to establish contact and educate fellow CCTLD managers as to the importance and severity of the threat imposed by Conficor.

With regards to Conficor, I think it is accurate to state that we fought and won the battle for control of the DNS. We CCTLD managers should be proud of our efforts in this regard. The worm has morphed into a peer to peer update control system that lets us rely on some rogue domain names have diminished considerably. So the state of the war against the worm is considerably more murky. It still infects millions of houses and there is little likelihood of that changing for the better in the immediate future.

But just because recent variances have shifted to a peer to peer control there is no reason for us as CCTLD managers to abandon our responsibilities to either sink hole or block names going forward. As we saw in the previous presentation we cannot abdicate our responsibilities to protect and defend the DNS against this and future attacks of this nature.

Partly as a result of the impact of the Conficor Worm on the DNS community the Incident Response working group was established in late 2009. This working group was originally chaired by Norm Richens and is now chaired by (1:34:59 name) and we will be brought up to date on the work of this committee shortly.

So at this point I'm pleased to acknowledge the participation of the fellow panel members. I'm going to start from left to right and begin with Ito.

Ito:

Thank you. Good afternoon and my name is Yuri Ito and I'm working ICANN security team. I would like to have a quick couple of minutes to introduce the development of the DNS concept and the statement and the status update to this workshop. So it's a global DNS CERT.

The background is the committee does instant response on Conficor and growing number of the domain hijackings and all these and also of protocol of vulnerabilities. There are also many things threatening and impacting to a larger DNS operator's community.

We just had a presentation from Dave on this new update. But these incidents really need global coordination as he mentioned. Those community mechanisms responding to those incidents, as a community and the need of the facilitator or maybe Incident Manager role in a community is needed. So the community calls for systemic DNS security planning and response and we hear that. Also we had this symposium, DNS Security Disability Resiliency Symposium which was the participants get together and listed all the risks and threats against DNS.

One of the risks that participants pointed out was the resource constrain organizations have difficulties to establishing networks or professionals to reach out and solicit information from. So having a dialogue with those resource constraint operators maybe in the African region or Southeast or Middle East or those regions we also find some of the gaps.

Also we did some existing response capacity analysis. There are a number of community efforts and framework that the key resources get together and develop remediation and work together.

We talked to them and started having this dialogue and also again identify that there is some gap to those key resources, to those resource constraint operators because some of those groups are membership, for example or maybe a very secret handshaking society that if you don't have face to face contact it is very difficult for them to be vetted and getting to the mailing list or community. So there are some gaps.

So when we're talking about the global and sustainable coordination mechanism to outreaching again to the global operators, we thought that maybe developing DNS CERT type of facilitation, Instant Manager role mechanism might be needed. So we started developing this concept of DNS CERT. and the mission that we put it under the public right now is the mission of the CERT is ensuring operators and supporting organizations how to the security coordination center with a sufficient expertise and resources to enable timely and efficient resource to threats to the security stability and resiliency of the DNS.

And the goal is again this mechanism to provide full time and global and sustainable coordination point to all those stakeholders, especially less resource constraint, less resource operators. To ensure a bridge to existing security organizations and the resources to those operators. And participation and feedback is most important to this program, so we are starting this dialogue with the regions TLD operators or DNS operators, also cyber security community and those existing response communities.

Also we're talking to the CERT community as well, like TF CERT or the regional CERT groups. Also, finding the needs of adding the capacity to those national organizational CERT as well. And so ensure resource constraint operators needs are met and keep collaborating and communicating with the community and find out their needs.

This document, the business case document is now on the public review on the ICANN's website. So please go and take a look at the business case and your comments, if you think it's a great idea or good idea, support it and then we'll move on to the implementation phase talking to the community and finding out what is a good way to implement this. So the document is up there and security and me and Greg my boss over there in this room and we're here for the week, so whenever you have comments or questions please grab us. Thank you.

Jorg Sweiger:

Hello my name is Jorg Sweiger and I'm with Dynek and currently here in my role as the newly appointed Chair of the Incidence Response working group. Actually I haven't prepared any slides because I was refusing to give a presentation because there is a presentation that is supposed to be given in the regular ccNSO meeting. I just basically said I don't have any new or anything spectacular that would give me the right to speak.

Nevertheless, I'm about to say a couple of words of what has been going on in the Incidence Response working group. As you have probably been realizing, the norm as the Chair of the

working group has resigned so we have been expecting some delays. After he resigned a new Chair was to be found and well that's basically why you are being faced or confronted with me right now and I hope you enjoy.

What we've been doing from that point in time onwards right now is that we just refocused our work and came up with a definition of what incidence response is all about and what should it really mean in the context of the ccNSO. So basically we haven't done anything other than setting up a completely new work plan that is operational. We defined what is supposed to be an incident for this working group and we will discuss and that probably is the issue that has been tackled here as well, how is the interaction between the Incidence Response working group on one hand and the planned DNS CERT on the other hand. We just didn't come up with a clear answer to this question, so this is going to be discussed.

Steve:

Thank you.

Roy:

Good afternoon again, my name is Roy Adams and I still work for Nominate even after my presentation from before. I was asked to sit on this panel. Now I don't have slides prepared either. But the concerns I have can easily be done without slides.

I come from a DNS CERT world in a previous life, sorry not a DNS CERT world I come from a CERT world in a previous life. I think and I need to introduce you to some vocabulary of what it entails to be a CERT. A CERT in general is like a voluntary fire department. It would be good to have a constituency and it would then be even better to have some jurisdiction over that constituency in order to do something as a CERT.

Now of course ICANN has a large constituency but it also doesn't have the jurisdiction over a large part of that constituency, namely the CCTLD's. So even though ICANN has basically a stick for the new GTLD's and the existing GTLD's what is the carrot for the CCTLD's? Next to that another concern I have is about the cost. The cost is what I understand projected to be about 4.2 million per year on ICANN's budget. Now that comes close to…it's not?

Ito:

It's not from the ICANN's budget.

Roy:

Oh it's not even in ICANN's budget yet.

Ito:

The number is calculated based on the same size and same service level of national CERT reference but it's not confirmed, of course, ICANN budget yet. It is the reference number.

Roy:

Okay so if this all goes through it goes on top of the existing budget? Okay.

(1:47:02 – inaudible no mikes)

Greg:

…who pays to be direct, so the bill is 4.2 million and the question is should ICANN take that bill, should it take a portion of that bill, should the community pay a portion of the bill? So Roy just to be specific we haven't determined the answer to that question.

Roy:

Okay thanks Greg that's clear. Okay. Then there are a lot of Registries out there that actually are DNS CERTS because they operate the national CERT itself. What I would like to see in this DNS CERT effort, I mean I actually like the effort behind a DNS CERT. I think there needs to be a lot more cooperation between CCTLD's and GTLD's, Registrars and their constituencies among mitigation of threats and mitigation of incidence.

The way it takes form with one organization doing this, I think you can use a somewhat smaller budget and go to existing organizations like the national CERTS and like, for instance, groups like DNS ORK and RASG and the likes and help them educate. Basically more like a starfish model than a spider model.

My last point is next to that there is already a lot of effort going on. We have at least I am part of several mailing lists that handles threat and incidences. Some of them which are unknown to the public and the reason for that is the secret handshake stuff but some of them are known like the NSD mailing list that basically helps organizations like Nominate to take down domain names.

Also for instance within DNS ORK there is a disclaimer here, I'm a Director of DNS ORK so I'm not here to start promoting DNS ORK but the organization exists and we do have meetings, we do have mailing lists, we do have threat mitigation and so like I said, I do like a major push forward for organizations like Nominate to join their efforts because DNS gets more and more abused on the internet. I'm not 100% sure about the module that is projected currently. Thank you.

Steve:

Thank you. As I noted at the outset, prior to September my focus on incidence response was cyber attack related which we've heard from our previous panelists, having been both a victim of denial service attacks and having to manage and deal with ongoing WHOIS server attacks.

However, early on the morning of September 29, 2009, the Samoan Archipelago experienced an 8.1 Richter Scale earthquake which was centered just south of Independent Samoa. Besides the immediate effects of the earthquake the seismic event also resulted in a tsunami that at its peak reached 14 meters in height as it approached the southern coast of both Independent Samoa and American Samoa and specifically into the harbor in Pango Pango.

With regards to the Registry, we lost our on island name server to the earthquake. The ISP where we were co-located suffered seismic but not tsunami damage. Seismic damage occurred to satellite dishes on the island effective satellite connectivity. Damage to the cellular telephone towers throughout the island was widespread in both the Samoa's. And the recently completed Samoa/Hawaii cable was also affected. So our immediate take away from this was several fold which I hope to detail at some point here further on.

More recently on the 27th of February Chile suffered a massive earthquake and subsequent tsunami and I will now introduce our well regarded colleague Patricio Pablito who will describe how his Registry managed the effects of their incident.

Patricio:

Thank you. Hello everyone I was supposed to be in Nairobi right now but something happened a little over a week ago that makes it advisable for me to say in Chile. As you all know, on the 27th of February in Chile we were all woken up in the middle of the night by a very strong quake. It seemed to go on for like forever but it was actually only 2 ½ minutes.

As you can see, you should be seeing a map that shows the location of the epicenter, which is south of Santiago and north of Concepcion. Concepcion is Chile's second largest city and which turned out to be the one that was most affected. The big red circle shows the magnitude of the quake and the smaller circles are the aftershocks. What we can see there are those that have been felt by February 28th. But since then we've have many of them, over 200 and actually there was one earlier today at 5 a.m. local time which was convenient for me because it woke me up so I could begin watching the presentations of the tech day.

In the next slide, you can see how big this quake was compared to other ones. It turned out it was the 5th largest in recorded history. That is in the top diagram you see at the right and there was another big Chilean earthquake before the Alaska one and that was in 1960. That holds the world record a 9.5. Now this one was 8.8 and one might think there isn't much difference between 8.8 and 9.5 but there is. The bottom diagram shows how much bigger the 1960 quake was compared to this one in terms of the released energy.

Now this one was very, very big on its own right. And if you compare it to the Haiti quake which was a recent one, it's about 500 times bigger than Haiti's quake. So we're really talking about a very big quake. And the next slide you can see a picture and that's become like a symbol of the destruction in Concepcion. You can see that tall building before pictured left and then you can

see the same building lying on its back on the ground after it collapsed. That shows how big the destruction was, how much of a challenge for everybody in Chile this earthquake has been.

On the other hand, the fact that we have this big earthquake back in 1960 implied that the building codes were very strict after that and I was talking to my colleagues in structure engineering and for them it's a real shame that this one building collapsed and taking several lives. It is actually the only building that collapsed in the whole country. There are many others that were significantly damaged and a lot of them will have to be demolished and that's a big loss of investment but not a big loss of life. People could walk out of those buildings but not out of this one.

So what kind of problem did that pose to us? We were in Santiago where the quake was very strong but we didn't have so much destruction as in Concepcion. I will talk about our non-DNS infrastructure and our DNS infrastructure.

The non-DNS is a website, the registration site, the back office processing, payment processing, email and all that. We have 3 sites for that and there is the main one and it's a very good data center. It is on power generator. We have another one in our offices that only one that has UPS and it a contingency site with UPS power generator. This electrical power is very important because at the time of a very big quake all power goes down. Actually you know the power of the quake is very strong when the power goes down. That is part of all the problems generated by the quake but also by design. Because turning the electricity off avoids some of the problems like fires, for instance, because of gas leaks.

So when there is such a big quake you can expect the power from the grid to be off. All production servers have 2 mirrors one on the same side and the other on the contingency side and all the network equipment is duplicated.

The next slide you can see the impact on this infrastructure. At the time of the quake, all the links degraded or most of them, ones were powered off, the orange link, and all the sites began operating with auxiliary power. Our engineering team was able to verify that the website was still operating a few minutes after the quake by using the Blackberry network. And about a half hour later, one of the links the Intel link went down and that was the one that was still operating at full capacity. And an hour after the quake we had people on the premises inspecting the offices which were in good shape and also the servers.

And the power went down in our office about 2 hours after the quake when the UPS went down. We don't have our own generator there. Now the main site is the one that holds small loads so we don't see that as a problem and also because our offices are located within an area surrounding the Presidential palace, sort of a protected area, that is even priority for the return of electricity.

At 9 in the morning more or less, the Intel link came back up and that signaled the beginning of the normalization of communication. Our team at 10 was at the main site for inspection and we had a problem with the orange link later and we had energy back in the office and the links started coming back up. The contingency site was inspected later. The zone generation was completed in the evening and at the end of the day everything was back to normal there.

Now the next slide shows what happened to our DNS infrastructure. We spent a lot of effort in trying to set up a very resilient network of DNS servers. We have 3 clouds; we also have one cluster and 2 UNIX servers. In Chile we have 6 servers active, 4 in Santiago, 1 in Concepcion and 1 in Balpariso. You may remember I said that Concepcion was hardest hit. And we also have a local mirror of the F root in our main site in Santiago.

Now what was the impact of the quake? Well all the international notes were unaffected as expected. So that immediately meant that for the world at large that field was still working without any problems. What happened in Chile, at the time of the quake we saw minimal traffic in some of our servers and that is explained by the fact that power went down in all the country more or less or at least a big chunk of the country where most of the population is located. And also many of the ISP's lost connectivity. So this situation of having minimal traffic lasted from 3:34 until 9 in the morning when we started seeing traffic again on our link.

The only server that actually went down was the one at our offices because of the USP running out. So by the next, at midday more or less we had a pretty normal situation. The server at Concepcion was down and it continued that way until the next day. We had very little information about what was going on there and the website for the hosting provider was down but the server went back up the next day.

So what kind of impact did we see on the traffic? On the next slide you can see what the traffic looked like before until 3:34, most of the traffic was…this is only the servers we operate. Most of the traffic was handled by the NS server and that situation changed. The server we operate in Los Angeles, California started taking most of the load and that continued for about an hour.

In the next slide you can see this also at different time scales. At the top left, you can see the same slot that you had in the previous slide and that's what happened in one hour. The one to the right shows 4 hours and you can see how big of a disruption this was and you can see that for 2 hours after the quake the situation was still very much irregular. Most of the load was taken by the site in Los Angeles and all the other servers had taken very little load.

At the bottom to the left you can see a full day and to its right you can see what a typical week looks like and again you can see the disruption in the traffic. But in spite of all that, we believed that at no time we were in a situation where we couldn't resolve names even within the country. Not that there was too much demand for that because most of the internet in Chile was down anyway. But we were able to resolve names at all times.

What have we've drawn by way of conclusion? We haven't had that much time for that yet, it's only been about a week and people have many pressing needs. We were lucky in that no one from our team was hurt by the quake nor had their houses damaged in any significant way. So we have been able to come with all our staff except for a few that were stranded abroad and unable to go back to Chile because the airport was closed.

But except for that we think we've faired pretty well. What are some of our conclusions? One is that the network of DNS servers around the world guaranteed uninterrupted domain name resolution service for .CL. that was our main goal and in all our plans for readiness for this kind of event that was our priority and it worked.

Having a local mirror of the root was important too because we haven't been able to confirm that but we believe there was a time, a certain time during this crisis where all the international links were not operating. So having a mirror of the root in Chile allowed for domain name resolution to continue during that time window. All the sites responded as expected.

The zone generation was normal but there were some problems that prevented it for being published at some times during the day. But at the end of the day that was completely normal. Not many changes were made to the zones during that day understandably although some people were registering domains as soon as we were able to take those registrations and a couple of domain names that had to do with earthquake, Earthquake2010.cl and similar names.

We are analyzing the events and we will do it to improve response for future emergencies. Chile is a country with big earthquakes so we should be prepared for something like this in the future. Hopefully not so big but there will be emergencies in the future. And one of our conclusions of this analysis, this very preliminary analysis is that one of the main problems we had was communication within the team. At the time of the quake, we still could call by cell phone for a little while or landlines but that didn't last for very long. For one thing because of the electricity going down everywhere including for the providers and also because of the congestion generated by everybody trying to call everybody else to see how they were after the quake.

So basically there was a time where we couldn't communicate. I was at home but except for a call that I could get through to the head of engineering division I couldn't communicate at all later. I knew people were going to the office and they really were going there but I had no way to communicate with them for a few hours.

So one of the things that we will be analyzing is setting up some more resilient means of communication, be it by radio or by satellite phone, something that will not depend on the local communication infrastructure operating. And this problem, by the way, happened not only to us but also to other people who should have ways of communicating like the emergency workers, relief efforts was all hampered. Even the announcing to the people that there could be a tsunami, there were many problems with that and one was communication.

So in retrospect, not only for us but for the whole country, we think that the communication infrastructure has turned out to be very fragile and we will have to do something about it. At least if we want to be able to guarantee that our team will be in communication within itself.

And in closing, I put the final slide I put this very strong photograph that has become very famous of this fellow who was trying to rescue something out of his house and the only thing he could find was this flag. It has become like a symbol of the will of the people to overcome this big thing that has been generated by the catastrophe that has happened in Chile. That is my presentation.

If you have any questions I will try to answer them.

Steve:

Thank you Patricio for that and I commend you guys for the resiliency and the design of your infrastructure and for how you guys bounced back as quick as you did. I think what you presented here is a brilliant case study for the rest of us to take away and emulate to a large extent. So thank you very much for that.

Patricio:

You're welcome.

Steve:

And my personal sympathy goes out to you guys having just gone through this myself. We have some questions coming your way.

Curtis:

Hi I'm Curtis Lindquist and I actually work for Net Know and would like to express my relief that all your team is unhurt. But with the query load, we actually noticed from our point of view when the earthquake happened that we every minute compared the server serials with the master server in Chile and we could no longer do that for a while when we noticed what was happening and news report coming in.

On our side, although we never had to do that, we quickly decided if we had lost all communication with Chile or had a longer time, we decided to not expire the servers and we would have kept the CL zones alive for as long as we could. In the end we never had to because connectivity was restored. But we did see a quite substantial increase of queries for that CL on our servers around the world that I suspect was the drop we saw on the graph that was in your presentation.

Patricio:

Yeah as I said we did generate zones but some of them were not published and that's also a problem that we will investigate because it may not only have just been because of the quake. But at the end of the day the situation was normalized. And thank you for taking the appropriate measures for the zone not expiring in case we were not in communication for longer time. I'm glad we could get back in communication very quickly.

Steve:

I think that's it on the questions. If there are none for any of the other panelists we do have a couple of questions from Jay Daily with regard to the cyber thing and I'm just going to hit on 2 of them. It is at the top of everyone's list I think and one is what barriers need to be overcome for incident response to work smoothly between Registries, Registrars and security companies? That one we'll answer today obviously.

Then the 2nd critical question that I think needs to be thought about is how can we as Registries trust the information we get from security companies?

I leave those with you as take aways because this is something the Incident Response Group needs to think about and we as Registries need to think about seriously as to how we're going to expedite the coordination and cooperation that is becoming more and more important as we go forward in these attacks.

Do we have comments or thoughts from our panelists on that, either of those questions?

Jay asked 2 questions, well he asked many and the ones I'm focusing on first is what barriers need to be overcome for incident response to work smoothly between Registries, Registrars and security companies? The second question is, how can we as Registries trust the information we get from security companies?

Ito:

How would you trust the information, it's really I have been working in a CERT community for 8 years and whenever we get the information it is your responsibility to really validate it. Again, when we release advisories we make sure that the constituency takes the advisory but also validates themselves and adopt or implement a remediation. The decision is really the constituency's decision. The constituency's responsibility. The system probably running on their own environment could be some proprietary environment or system. It is really the information and we try to…I mean it think security companies or security community releases the information as accurate as possible and indentify the risks and communicate in an accurate manner. But it's I think the recipient or constituency also has the responsibility to how to respond to that. It's your analysis, your validation and your responsibility.

Dave:

This is Dave Dagan and I wonder if I might comment on that. I know I was on the previous panel but…

Steve:

Certainly.

Dave:

So currently there are several groups that work closely with Registrars to seek mediation of malicious domains. The general pattern that is followed is that a researcher will use some messaging system, generally email these days, to complain about a domain. And they'll also provide some proof, some indication in the email that what they saw was malicious.

The idea is that anyone reading the email could replicate or recreate the proof themselves perhaps by clicking on a link or reading through some of the trace. This can become difficult if this becomes the minimum standard because often times malicious sites take themselves offline for a period of time, perhaps they give inconsistent answers and the reproducibility of malicious results is not always guaranteed but it's currently the practice that's being followed.

I have seen a couple of exceptions to this generally around the area where the person complaining is from a government agency and in those cases often times the people receiving the complaint just act on it without asking for further proof. It's simply a function of the fact that they have built up a rapport or reputation.

I say this and I fully acknowledge that I don't think that scales in an automated fashion and it doesn't account for the fact that people may change jobs and professions and all these relationships now need to be rebuilt. In a sense, the system we have now is fairly ad hoc.

I think it's beneficial if the Registries could begin to trust certain operators within those communities, for example, a CERT that could validate or verify reports would be one level of proof. If they end up trusting some of the large research organizations that are frankly never going to go away and are well known the Semantics, the MacAfee's the large groups that are consistently staffing this sort of concern that would be useful.

At the far end of the spectrum is this need for judicial proof, some "order". And I understand why that is a standard and why that standard is applied but I wonder if we can't come up with something short of that, at least for myself I'm willing to say I'll get insurance and I will stake my reputation on the reports that I generate. Perhaps not every researcher is capable of doing that but it is something I would offer as an intermediate. Because the current system of requiring strict judicial proof is I think a bit balanced too far and often proves to be an inconvenience that Bot masters are increasingly aware of.

Steve:

Thank you. We have one closing comment from a panelist and then we're going to have to call this session to an end. We will conclude with some concluding Tech Day remarks for Dr. Eberheart.

Roy:

I have thought about answers to the questions from Jay. How can we as organizations within the DNS space work better and more closely and get rid of the barriers between us? First of all, what I think is appropriate is that an organization will have an entry point for these kinds of issues. For instance, a responsible person or a mailing address within a company that people can send mail to, such as an abuse contact.

Then it would be good, for instance, a lot of Registries, I know a lot of Registries are already highly active in this area of threat mitigation. They are visible on mailing lists, they talk to each other, instant messaging, they see each other at meetings and yet they just don't call that a CERT even if in fact it is a kind of CERT organization.

So it is simple step to organize that and subsequently register yourself at first, the first is the form of incident response security teams. For instance, there are 17 or 18 organizations in the UK that have registered as a CERT. this is not rocket science. This is fairly trivial to do. That immediately follows by some kind of trust building. You don't have instant trust. If you work with security companies and you get information from them, of course, you need to validate that information. But if you cooperate in organizations such as FIRST and assert yourself then others can vouch for organizations as well. This is how you build trust between organizations.

Typically, in these kinds of organizations the technical people will go to these meetings, the technical people who understand the technical situations involved in this. For instance, the folks you will generally see here at these ccNSO meetings.

I think with that I tried to answer Jay's questions.

Steve:

Thank you.

Male:

May I add something to what Roy just said? If I got right what Roy just said, what he depicted is that we should provide an infrastructure for communication and for reaching each other. So that is exactly what the Incidence Response working group is all about. We are up to building a contact repository; we are up to build this repository with information about communication channels. We are going to describe whether or not we would need secure means to communicate with each other and so there is something in place right now that is going to tackle at least this part of the problem.

Steve:

Thank you. I want to thank the panelists and I want to thank especially Patricio for his presentation from afar. That is it for Incidence Response. We have a lot of work to do and if you're interested in the working group by all means continue to follow it. Thank you.

Chair:

Thank you very much. We are supposed to be out of here because the room is needed for another working group. So I call on Norm to wrap it up and he didn't really want to wrap it up, he wanted to only give comments about how the remote participation works. I'll abuse him and ask him to wrap it up as usual anyway.

Norm:

You got me. First of all, I've been working with this group for several years now and if you don't know I'm no longer with CRM and I'm with ISE now. But as I started to work with Eberheart in setting up this session I continued to do so.

I'm just going to say a few things and I won't talk about the remote participation other than to say I'm very happy that it is there and I hope ICANN staff continues to improve it and make it available across all sessions.

On the cyber security side of things the way the world has changed now and whether you're in the domain business, DNS service provider or in cyber security, those 3 areas are verging. It doesn't really matter do you think you should get into it or not? Well I think everybody is. We really have to just decide on what policies you want to put around that.

As I said before, the good, the bad of CCTLD is that you make your own policy and you implement those policies and then you enforce them. But I think it's very, very important that the CCTLD's work together on this because if you don't others will do it for you.

Chair:

Thank you very much. Thank you very much everybody for participating locally and especially the ones who used the remote. We had a little bit of hiccups but that's the way these things go. I am quite sure that next time it will be better. I'm quite sure from my understanding it was quite understandable and my son in Germany sent me a few messages. He was listening and he could understand very well, so I reckon the connectivity is good enough. We'll work on it so that on the next ICANN meeting we can do it again.

Enjoy your time in Africa.