

ccNSO-members-2-09mar10-en.mp3
01:07:38
Transcribed 03/25/2010 LNC

Chris Disspain: Good afternoon ladies and gentlemen, welcome back from the GAC.

Francisco, are you here? Is he? Oh, okay.

Gabby, is Francisco online or who's Chairing?

Me?

Francisco Arias: Can you hear me?

Chris Disspain: Yeah, is that you Francisco?

Francisco Arias: Yes.

Chris Disspain: Hi, it's Chris, how are you?

Francisco Arias: Fine and you?

Chris Disspain: Very well, thank you.

We've just come back into the room from the GAC session. Give us a couple of minutes and then we'll start if that's all right?

Francisco Arias: Okay, sure.

Chris Disspain: Okay, ladies and gentlemen, we're going to move into our next session which is the IETF Data Escrow Initiative. And that's going to be a presentation by Francisco Arias for the gTLD registry, technical liaison. Francisco, I have the slide changer in my hand so if you just tell me when you want to move on to the next slide, I'll be happy to arrange that for you.

Francisco Arias: Okay.

Chris Disspain: And please start when you're ready.

Francisco Arias: Okay, thank you Chris.

Hi everyone. This will be a short presentation.

This is – Chris could you go to the next slide please?

Okay. So Registry Data Escrow is something that has been done for the gTLD for five years and is aimed at ensuring registry operations continuity.

For the new gTLDs process there was new development, we were developing a new version of this Registry Data Escrow specification and in the Seoul meeting, people from the gTLD community suggested us to move this Data Escrow work to the IETF. We also consulted with other members of the community, including some people from ccTLDs and we were given positive feedback about this idea. So we decided to move on with this.

Turn to the next slide please.

So this effort focuses exclusively on the technical point of view of Data Escrow. The idea in the IETF is not to discuss anything besides the technical side of the Data Escrow.

The goal is to produce a specification that allows a registry B to rebuild the services previously offered by a registry A – another registry – with a minimal impact to registrars, registrants and end-users.

And we think this move to the IETF makes sense because the IETF is technical forum where technical people normally feel more comfortable working with and we think at the end we could have many more people participating here, technical people, and so at the end we will have better specification as an end result.

Next slide please.

What is NOT the Data Escrow effort in the IETF? This is NOT an invitation to do Data Escrow, but for those that already care about Data Escrow, we know about this do already the Data Escrow for themselves, at least (inaudible 00:20:31), as far as I know. And we know that some others may be interested for different reasons. So if you care about Data Escrow, please join the support in the IETF.

And also another clarification, there is no intention whatsoever to delay the new gTLD program waiting for this spec to be finished. We know the IETF and ICANN are different organizations and we cannot depend on the work of the other so if the specification is ready by the time it's needed in the new gTLD process in ICANN then it will be used. If not, they will have to see what to do, probably use the advice in the Application Guidebook Version 3 or maybe use the version that most advanced version at that time inside the IETF. We will have to figure that out at the time that it's needed.

Next slide please.

So the current status of this work is we already have a discussion list. There is an active mailing list. We have had several register. There will be a BOF in the IETF meeting in Anaheim in two weeks from now. The BOF, by the way, is the name in case you haven't heard this term, is the name that the IETF gives to our working group information. Our group that is still in the first stages of formation.

This BOF has two experienced IETF co-chairs: Edward Lewis from Neustar who probably people know because he was one of the co-chairs of the ETP – the IETF working group that developed ETP – also the other co-chair is James Galvin, also an experienced IETF person who has been involved in many DNS working groups inside the IETF.

At this moment the discussion in the mailing list is focused on scoping of the work.

Please, next slide.

On the mailing list there are already people from gTLDs and ccTLDs already participating. Some from Afilias, we've got BG, CH, France, JP, Neustar, Netherlands, UK, South Africa and probably are more others subscribed. I know of some others that aren't listed there.

So we expect to have formal consideration for this group to be a formal Working Group in the IETF. We expect this to happen after the Anaheim meeting.

Please, next slide.

So, if you are interested, this is the link to the mailing list and you can subscribe there. You can see also the mailing list archive.

And as I said, there will be a meeting for this BOF in two weeks on Wednesday, March 24th from 10:30 to 11:30. So if you care about this topic, join, you are more than welcome.

I guess at this point if you have any questions, I'm here. Please let me know.

Chris Disspain: Thank you Francisco. Are there any questions or comments in the room for Francisco?

Okay, no Francisco. Thank you very, very much indeed.

Francisco Arias: Thank you.

Chris Disspain: Thank you. See you.

Francisco Arias: Bye.

Chris Disspain: Okay. So the next session is the Incident Response Plan Working Group, Jorge are you? Here he comes.

And that will be followed at 4:30 by the Delegation/Re-delegation Working Group report and discussion. And then the close of the meeting for the day.

So are your slides already over there?

So we should have some slides, that would be great? There's the changer. You can sit here on the television.

Jorge Schweiger: So here, yeah. Good day. My name is Jorge Schweiger. I am the CTO of .DE and I'm chairing the Incident Response Working Group as well. So it's a pleasure to give you an update on the work in progress that has been made by the Incident Response Working Group.

First, a short overview, what the purpose of the group really is. And that is here for a reason because we have been discussing that. So purpose is to implement mechanisms for the engagement of and the interaction with ccTLD registries during an incident. And such incidents for sure that do impact the DNS.

So to be more precise, the scope of the Incident Response Working Group would be to set up a repository of ccTLD contacts and the corresponding channels of communication for incident response. And we are tasked to qualify what really is considered to be an incident and we are supposed to describe some escalation procedures and action paths.

Sorry, I should have glanced at the slides here.

So let's move on to the next one.

Well, according to the charter, this is the work plan we came up with and it consists of four points.

First of all we want to define what is considered to be an incident according to this Working Group. And we want to do that actually until March 10th. So there has been going on a lot of retirement on that issue and we are pretty sure that we already had defined what is considered to be an incident quite good. And basically I'm just waiting for your input to make it a standing definition.

Second point on the agenda would be define use cases of how this contact repository is supposed to be used. And we want to specify that until April 30th.

Next thing would be to define escalation procedures and action paths. We want to do that by May 30th.

And after we've done, especially the latter two points, then we would be in a position to specify how the data model of this repository – what it's going to look like – so what we're heading for is that towards the Brussels meeting, we'll have defined the data model for the repository and we'd know exactly how it's going to look.

And finally, we would like to come up with suggestions on who actually will implement this repository. We want to specify and suggest who is running it, who is maintaining it and if we do know all of this, the second question would be at what expenditure this contact repository is supposed to be run.

And I think that we are capable of suggesting a pretty good draft until a month later and then Brussels.

All right then let's get to a little bit of work we already had done.

So the definition of an incident we came up with and for sure this has got to be revised frequently, is as follows. So an incident is a systematic, rigorous preparation of or actual attack on; first, the availability of the DNS and registration systems, attack on the data integrity or privacy of the DNS or registration systems, attacks on the stability and security of the Internet at large, and that is where a coordinated international response by operators or supporting organizations is advised.

On the right hand side of this slide you see some attack vectors that refer to the bullet points on the left just to give you an example. So for example for sure, the availability of the DNS is being attacked by a DDOS attack and, as we are all or most of us are using EPP as a protocol, so a flaw in the EPP protocol might be exploited towards the availability of a registration system.

And what I do think is really interesting to see that we are not only focusing on the DNS but we did include registration systems as well.

Another point I will bring to your attention is that actually we started out with the definition giving just first two bullet points and then we've been discussing the third one as well, meaning this one, that. Because we felt that the first bullet points would not even cover Conficker and as this Working Group kind of stems from the Conficker incident we felt that we are in need to

enlarge the definition a bit to cover any attack that might happen in the future. We don't even know what it's going to look like right now.

So having said that, it is clear that the following is not to be considered an incident for the purpose of this Working Group. And that is the malicious use of the Internet itself, so we are not talking about SPAM distribution, we are not concerned with the unlawful use or misuse of specific domains referring to content, and it's not our business if any routing problems occur with BGP or something else.

Okay, this is as far as the results. There are some questions open as well and I would really appreciate if the audience could comment on each of those questions because it would really be lovely to get some guidance on that.

So first thing would be that we weren't sure or are not sure by now about the sophistication level of this contact repository. So, first of all, who is entitled to access? Are we in need of certain security measures and communication means that are really secure to access those repository? And finally, the most important question, who is going to host this repository? Is this supposed to be hosted by a professional third party like TI for example? Whereas TI, to those of you who are not familiar with the term, is a company that is already hosting repository for the CERTs and that may very well be adapted according to the needs of the contact repository for the incident response working group. Is this supposed to be run by a dedicated ICANN branch? And I'm thinking of the DNS CERT that has been under discussion after the last at least couple of days. Or, at the lower end of the possibilities, may just be a simple IANA database extension?

So are there any comments on that?

Well, if not. Robert.

Robert Shishko: My name is Robert Shishko, I am Technical Manager of NICat but I also happen to be the head of the Austrian CERT so therefore I know some of these procedures very well. May I comment on this?

I am very much in favor of having a very lean database structure there. (Inaudible 00:35:44) introduces an excellent example of running an excellent infrastructure at actually no cost. And I think it's not right that this really a company, it's more a joint effort of the teams there. I think an amendment to the existing IANA database would be fine enough. And regarding to the security issues, keeping the data confident or not, it's not so much about the confidence of the data, it's more about what kind of data do you really to publish and do you want to publish your secret emergency number and be waked up in the middle of the night. So that's all about it. So I think you could work with two levels of security, one for the public and one for

emergency contacts because you really want to keep secret in a way that is not triggered an atomic alarm in the middle of the night just with someone who is not happy with your service calling you.

And one remark on this context you mentioned DNS CERT. What really puzzles me a little bit is that we now have an Incident Response Working Group and you might know that CERT stands for Security Emergency Response Team – sometimes CIRT in equivalent terms stands for Security Incident Response Team. So the DNS CERT or C-CERT or whatever the name it might be, seems to me a generic part of this Incident Working Group or the other way around. And while I think everybody seems to be in favor initiatives helping to improve cooperations and response to security threat and therefore I suggest that at least the Incident Response Working Group should take the task of collecting the various feedbacks from the community. I think everybody seems to be in favor to improve this cooperation but the meaning of having it – what is inside the scope, what is outside the scope – how broad should this organization be, be set up very lean or having really tasked on hands on experience or do they only focus on communication. That really depends to whom to you talk. And I would really suggest that this Working Group make take the cumbersome approach to collect all this input from the community and produce a kind report on what kind of, what does the community request and what kind of service this CERT should deliver to the community.

Chris Disspain: Thank you. Can I just suggest purely on a logistical and on a process point, that's a very interesting point. What I think would need to happen is that the Working Group would need to come to the Council – if the Working Group accepts that point and decides that it wants to do something – it then needs to come to the Council with a proposal and then we would consider that and if we agreed we would amend the Working Group charter. I think that's a very interesting suggestion so if you guys want to kind of work on it and come back to us, that would be, that would be excellent.

Robert Shishko: I just think we should not end up with two parallel Working Groups having the same topic in their name. That doesn't look very professional.

Jorg Schweiger: Thanks a lot and I'm really going to follow up on that one.

Yuri Demchenko: Hi, Yuri Demchenko. Actually – I am not so tall – okay, actually this is a bit a repeat of comments and discussion at the security, ICANN security meeting was on Monday. That was discussed again DNS CERT organization project. This major issue is CERT and Security Response is a quite well developed procedure and infrastructure in Europe and actually (inaudible 00:39:25) work only for Europe. It's European project. So if you want to cover the whole world, you need to extend (inaudible 00:39:37).

And next, your definition of incident is quite good shape for the DNS but it probably also should be aligned to the general definition of incident. Like NRFC dictionary or special couple of RFC that published for that incident, object description format. But it's also incident object description format that implemented in particular by C-CERT in the United States and mainly CERT in Europe. So generally when it's a possibility to provide you more information about this. So this but okay...the question is not quite good, this only a comment.

Chris Disspain: Yes, thank you.

Unknown male: Yeah, Jorg, question from me. I think I saw in your definition of an incident that it explicitly stated that this was about attacks. And it seemed to include incidents as a result of human errors for instance? And I ask that question because I think also there this kind of a structure could help. In our recent Centa meeting, we had a presentation by one of the European registries who sometime ago published a corrupt zone file. And in order to minimize the consequences of that action, they had to get in touch with all kinds of parties in other countries. And they didn't have their phone number or fax numbers or telephone numbers. So I think that is also some kind of an incident which is not the consequence of an attack but where a structure like this one could help.

Jorg Schweiger: Agree.

Unknown male: Can I read this definition of incident from RFC 3067? "An incident..."

Unknown male: The line is closed by the way.

Unknown male: "...is an security event." The essential issue is that event but this event may be either one incident or serial of incident. So maybe it will also work to...

Chris Disspain: Okay.

Simon McCalla: Hi Jorg. Simon McCalla, Nominet U.K. A quick question, with regards to the fact that there's all these Working Groups that are currently trying to define similar grounds of this, is there work that the Incident Response Group could be doing now that works in the informal structures that we already have in terms of mitigating attacks, in terms of just having circulating contact details and having a really immediate impact?

Jorg Schweiger: I'm not sure whether I got your question completely right? If you are referring to in which way are we interacting with the groups that are really in place right now, to be honest, this is actually the third point on the last slide. There's clearly and I do feel the same, there's clearly a need that we do have to consult, that we do have to relate and that we do have to clear on any or

with any of those groups who are currently around. And if you would like to take a look at this slide, you will see that this slide is providing what I think is a pretty comprehensive list. So the answer is, we are on our way to do so.

Simon McCalla: Okay, thanks.

Unknown male: Yeah, I hope it works, yeah. The Incident Response Group met on Saturday and one of the outcomes was that the ccNSO website will include a link to all these different types of organizations so that at least it is clear which organizations exist or are currently in existence and as Yuri said, it is not to overdo or do work of others.

Chris Disspain: Thank you Jorg. Could you join me in thanking Jorg for his presentation?

(applause)

Jorg Schweiger: Thank you.

Chris Disspain: Thank you very much. Okay, so, Keith, Becky, if I could get you guys up? The last session is the – today – because I don't what it's like down there but I can tell you it's very hot up here – the last session today is the Delegation/Re-delegation Working Group.

You've all had a copy of the progress report I hope and many of you, some of you were actually at the workshop on Sunday when we talked about experiences. But the Delegation/Re-delegation Working Group is chaired by Keith and Becky is the Vice-Chair. So over to them to take us through what they need to take us through.

Keith Davidson: Good afternoon everybody and that's really great to be the last session on a warm afternoon. So I think those who want to have a snooze, they probably had better leave the room, because this is going to be a fairly exciting session.

Can I or Gabby, oh thank you Gabby, you're already there.

Can I take the liberty of just running through the slides that we presented at the workshop the other day and can I also apologize to anyone who's been through this discussion just before in the GAC and at the workshop itself? Sorry for the repetition.

These were the slides for the workshop rather than for this so I'll just run through the important points.

The workshop that we ran, just incidentally, was to assist the Working Group in understanding issues and experiences of others relating to delegations and re-delegations that had occurred. Clearly not though to get involved in any

individual case of a delegation or re-delegation or re-litigate any earlier decision and our work is focused on overarching policies and procedures.

Next slide, Gabby?

Now, just to cover the Working Group. It is a Working Group established by the ccNSO Council. It's tasked to compare policies and guidelines with IANA procedures and discuss issues with ccNSO members. It will report finally back to the ccNSO on any aspects that may require a PDP.

Next slide.

The Working Group will report back to the ccNSO on three separate aspects over the next year. We decided to start with what we thought was the easiest, delegations, proceeding through retirements, and then finally re-delegations which we thought was going to be the hardest part. However, there has been overlapping discussions and the discussions the Working Group's had to date and it's not possible to fully isolate them into separate streams of work. And we're already behind schedule as the Delegations Report was due at this meeting and now hopefully in Brussels.

Next slide.

Part of the work that we've done – we commenced a review of the policies and guidelines and particularly RFC 1591, ICP-1 and the GAC Principles. And within the Working Group, I think it would be fair to say that there's no perceived significant difference between these individual policies. We matrixed them and looked at each of the policy statements line by line and while there are slight differences, there hasn't been anything of significant concern to anyone in the Working Group.

Next slide.

And the Working Group is currently researching ICANN Board decisions on delegations, retirements and re-delegations and our next set of work will be to compare those decisions and their alignment with the policies and guidelines.

And we are meeting iteratively. We meet up to twice a month and we never make a decision at a single meeting. It must go through two iterations due to the high numbers of people on the Working Group to give everybody an opportunity to participate in a reasonably friendly time.

And we've had a Working Group meeting here in Nairobi and also the subject did come up in the GAC session for some discussion, probably the biggest issue for those who were not on the GAC was the GAC having some concern that the ICANN Board perhaps didn't think that the GAC Principles on the

delegation and re-delegations of country codes was formally received by the Board.

And there's a link to the Working Group documents. There has been a progress report published.

I think that ends my presentation and Becky Burr provided a very useful background on the IANA function and the history of IANA's development to the workshop so I've invited her to give that presentation again, just to help round out the knowledge. And after Becky's presentation, we really just want to open this up for open discussion. Thank you.

Becky Burr: I'm going to apologize for a busy slide. I know you're not supposed to put this many words on one slide but most of it just background information.

In the Working Group we've been discussing sort of the origins of the policy or the practices regarding ccTLD delegation so I thought it might be useful for all of us to reflect on that.

In the 1960's the U.S. Defense Advanced Research Project Agency, known as DARPA, funded first UCLA and the later the University of Southern California where Jon Postel was doing post-graduate work, to maintain a list of host names and addresses and eventually these functions came to include responsibility for parameter allocations, IP addresses and the like, and over time became known, the functions became known as IANA.

There was never a standalone entity or organization.

In 1983 Paul Mockapetris and Jon Postel developed the DNS. He delegated .US – Postel delegated .US, .UK and .IL in 1985. And then between 1985 and 1993 others were delegated on a first-come, first-served basis and largely to people that Jon Postel personally knew in universities around the world but in cases where there were not universities, he delegated them based on his views about who would provide service to the relevant community.

In 1994 he published RFC 1591 to articulate, describe the manner in which ccTLDs were delegated and the principles on which they were delegated and the expectations for ccTLD managers and indeed, TLD managers.

In 1998, following recognition of ICANN by the U.S. Department of Commerce, the IANA functions that were at that point still being provided, undertaken, by the University of Southern California were intended to be transferred to ICANN. Because USC was working and developed various intellectual property under a U.S. Government procurement, they had the need for some certain formalities in transferring that intellectual property which is why the IANA functions contract came into existence. It really was never

intended to be sort of two separate functions. To be clear, the activities are quite intimately related in our view. Or in my view anyway.

Next slide Gabby.

So RFC 1591 principles, which is really, you know, whether it's a policy or anything else, it's the authoritative text under which we live and has a special place at least in my heart and I think in others. The principles were fairly straightforward.

Jon's view was managers had to be fair, just, honest and competent. That managers of cc's would be trustees for the delegated domain with a duty to serve the community. He specifically said concerns about rights and ownership are inappropriate. ccTLD managers should be concerned about their responsibilities in serving the community. There was a provision that required managers to be equitable to groups that were permitted to register in the domain, not that everybody had to be permitted but that whatever the rules were, they had to be administered fairly. He said that significantly interested parties in the domain should agree that the designated manager is the appropriate party and the government of the relevant country or territory was a significantly interested party among others. And the basic requirement was that the designated manager was required to do a satisfactory job of operating the DNS service for the domain.

Gabby, next slide.

The other thing that I think it's really important to keep in mind is that Jon had a list of things that were driving him when he put 1591 together and although this slide is largely derived from an RFC that was drafted by Jon Clensen some years later, I spoke with Jon about it and was impressed by the wisdom of his, you know, here are the death trap things to avoid. His goal was to never be involved with determining which entities were countries and which ones were not. And so he picked the ISO 3166 list. It may not have been a perfect list but it was, he just wanted an authoritative text. He did not want to be involved in determining who was or was not the legitimate government of a country. And he didn't want to be involved in prioritizing rights within the government. He wanted, in fact, he wanted disputes regarding delegation and operation of ccTLDs to be resolved in country and that is why there was a requirement that all ccTLD delegations had at least one manager who resided in the country. I know that there have been some exceptions in the margins there. But his view was that if there was somebody subject to the jurisdiction of the relevant government or territory, then the relevant government or territory had the rights and authority under whatever rule of law prevailed in that government or territory, to reach a solution with the ccTLD operator. So when disputes arose, Jon's direction was essentially to tell the parties to the

dispute to go away and solve the problem and come back when they had resolved it.

So that's the, in terms of documentation, that is our starting point. We are looking to document the history of the delegations and to the extent that any of you out there have history, have correspondence with Jon, we would love to have it because we would like to compile a sort of historical record, both for purposes of doing our work and also for purposes of history.

Keith Davidson: Thank you Becky. And I think when Becky refers to old documentation; I think probably mostly pre-1990. So if any of you have any records related to a delegation or a re-delegation prior to 1990, even if it's just email exchanges, we would very much like to hear from you. And perhaps not necessarily here and now but if you contact Becky or myself after the session.

I think in terms of RFC 1591 – thank you Becky, that's a very good clarification – and of course 1591 refers to more than just the operation of ccTLDs – but as the principles that you've raised that relate to ccTLDs go, I'd just like people to think about it for a few minutes and if you have any issues, disputes or items that you don't feel comfortable with in RFC 1591, to just come to the microphone and, you know, let us know what you perceive as a problem with that document.

Other than that, I think Chris do you want to provoke the audience into some participation before they all go to sleep?

Chris Disspain: I have a feeling that this audience is beyond provocation.

I think that we, I mean, we've basically delivered the message from the Working Group. We've asked for any documents that people have. I don't really think there's anything else.

If anybody wants to ask a question. Oh, okay, go ahead.

Unknown male: I just want to – one observation – first question.

How is Working Group consider the GAC Principles? Because from my experience and knowledge, this is most structured approach and it clarifies something that is not well defined in the PCP and RFC.

Chris Disspain: So the question, the key question there is not really how the Working Group perceives the GAC Principles, but more how the GAC Principles are perceived by perhaps the Board of ICANN and the GAC themselves because that feeds into a discussion about what documents are actually out there.

The GAC Principles, the 2006 GAC Principles are an official GAC advice to the Board and that document has a certain weight in the ICANN By-Laws. And it's, I mean, that's basically a fact, so that's the way it is.

The GAC Principles are actually often quite useful document because they are written in a way that enables argument about certain points. But they are, my understanding and I think the understanding of the Working Group, is that the GAC Principles are official advice to the Board.

Unknown male: Okay. But (inaudible 01:01:27) at least a known to me case of re-delegation still ongoing more than 10 years for .UA domain, Ukraine.

Chris Disspain: Well, okay. So sorry, I'm going to stop you. We won't, will not, discuss individual delegations because it's just not appropriate to do so in, you know, in a sort of public environment like this. So if you want to talk about something specific to do with a re-delegation, then you really need to not do that here. Okay?

Unknown male: I'm not talking about re-delegation.

Chris Disspain: Okay.

Unknown male: I'm not asking that.

Chris Disspain: Okay, so go ahead.

Unknown male: I'm not asking for a decision or an opinion. I simply know that there is a big gap between practice that is provided by or conducted by IANA according to PCP and RFC and GAC Principles. So for this case particularly, GAC Principles provide the most structured approach of defining relation of government to the, defining who is the organization, and this should be also in this document at least as a case where we have a kind of number of problems. Why is process last for 10 years? I'm not talking why, because I know some kind of as Chris asked, but still this should be reflected in the review or recommendation of the policies.

Chris Disspain: Yeah, no, thank you for that. This was the very subject of the workshop that we ran the other day and so that's rather unfortunate you didn't see the information on the ccTLD lists and so on because that would have been a perfect opportunity for you to make a presentation. But then the Working Group will take that into account and we may run another workshop at a future ICANN meeting.

I think information about an individual delegation or re-delegation is useful to the Working Group. Not that it can act on it, it's out of our scope to act, but it is useful information, particularly if you can provide definitive examples of

what you believe to be a gap between the policies and the practice. So you know, that can help our overall work, but it can't help your individual instance.

Unknown male: Yes. Surely. We are not pretending to want.

Chris Disspain: Yes, I understand.

Unknown male: So if I will have opportunity to provide individual way or to the mailing list, this will be, best effort to instruct.

Chris Disspain: The quickest way to do it right now is simply to provide it to Keith. If you do that, then we can take it into the Working Group and go from there, all right?

That's the quickest way.

Unknown male: Yes.

Chris Disspain: Okay, thank you.

Anything else? Are we done?

Keith Davidson: Just a clarification back from Bernie Turcotte who's our contracted ICANN staff support for this Working Group – I said we're looking for old documentation pre-1990 – he's quite happy to take any documentation pre-2000 but with a particular focus on pre-1990. So keep that in mind if you have a good email archive.

If there's no other questions. Oh, are there any questions or observations relating to the Working Group in general and where it's heading?

No.

Chris Disspain: Okay, thank you Keith. Thank you Becky.

Two bits of logistics. Tickets, if you are registered for the ccTLD dinner, if you are registered already, if you're not, sorry too late, if you're registered and you haven't got a ticket, you need to see Gabby.

Secondly, Gabby in her usual way and efficient way has published a survey for us to fill in. It's best if you do it today because otherwise you'll forget and then again tomorrow for tomorrow.

And finally, we start again at 9 o'clock tomorrow morning. The schedule originally said 9:15, we've changed it back to 9 o'clock. 9 o'clock in the Shimba Hills room which is on the first floor of this building.

And other than that, I'd like to thank everybody for their attention today. I'd like to thank the brass band for providing us with the music. And I hope to see you all again later.

Thank you.

(applause)