TORONTO – Working Session: Sunrise and Trademark Claims Implementation
Monday, October 15, 2012 – 17:00 to 18:30
ICANN - Toronto, Canada

KAREN LENTZ: Thank you. Welcome to the Trademark Clearinghouse Working Session on the Sunrise and Trademark Claims Implementation. This is going to be, as advertised in the title, a kind of nuts and bolts meeting. Going into some of the details on how we implement the new processes that were developed as part of the New gTLD Program associated with the Trademark Clearinghouse.

There is a session on Wednesday as well on the Trademark Clearinghouse which is a more general project update session in which you'll actually get to see the Clearinghouse in action. But what we want to do here is have a discussion and this room is not ideal for that.

So if people would be willing to move to the center of the room that would be helpful. There are mics on either side so it's intended to be interactive and to have a discussion with the participants here.

So just a couple of words about the Clearinghouse and the objectives that I think we share. You know that this is something new. It's something that was proposed with the idea that having a reliable source, an efficient source, with verified trademark data would create value by providing efficiencies in the process by reducing costs and creating a really useful tool.

I think that's where we're all trying to get to with this piece that we're working on now in terms of the registry interface. I'm going to just give a brief introduction then get out of the way so we can discuss. So in

terms of what has happened since the last ICANN meeting in Prague, we did a couple of things.

One was start a technical mailing list for potential new gTLD registry operators, those who would be responsible for the technical implementation of these processes on the registry side, had some requests for more interaction on that front.

So we created that mailing list about the same time we made some revisions to the draft implementation model that had been produced in April, also based on feedback that we got during the meeting in Prague.

As the discussions on the mailing list continued, there were some additional options and permutations to these processes circulated. We had a forum in Brussels where a number of people attended, hosted by Deloitte.

And a number of people attended remotely as well going through in pretty great detail what the concerns of some of the registries were. What the goals were of these processes, and how we thought that we could consider those as we moved forward.

There was a group of a few people participated in that who drafted some alternative proposals that were completed I believe in September, about a month ago.

And, you know, some of it new work since the Brussels discussion and so those have been something that we've continued to discuss. Those discussions have happened for the most part with a group of registry operators or registry focused people.

But a lot of the things that were discussed we recognized affect many stakeholders. These are processes that affect rights holders that affect registrars that affect ordinary individual registrants registering domain names.

So we wanted to make sure that we were balancing and making sure that we had a discussion about the "technical" aspects with a broader group so that we could raise some of the implications of the technical questions. So we have a few issues that we wanted to go through.

We're going to start with the Sunrise process which comes first in the chronology of startups so the Sunrise is, I think probably most people know. But it's a stage where there's an opportunity for eligible rights holders to register domain names in a new TLD before it's opened up for general registration.

We also have in regard to the claims process some discussion about encryption, some discussion about what the timing is for particularly in different types of registry models.

And then hopefully we'll get to some conclusions and next steps with some other issues being brought in along the way. So we'll start with Sunrise. Is Chris right here?

All right, so Chris is one of the authors of the Alternative Sunrise Model that was drafted. He's got a set of slides here that he will go through some of this for discussion.

CHRIS WRIGHT: Sorry, just one second. Okay, so firstly I'd just like to let you guys know that this presentation was put together in about a half an hour about two hours ago. So I'm sure there's lots of mistakes and so forth in it.

But I'll do my best to talk through them. I was asked by Karen and Kurt to step you through what the registry/registrar community proposed Sunrise model was and how it worked. Because we're using PKI and apparently there's a bit of confusion out there as to what the PKI actually is and how it works.

So I'm hopefully going to simplify it enough so that we understand what the proposal is. And the rest of us can make a decision whether we think this is a good way or not. I just have to keep my laptop and the screen there in synch.

What's the goal for Sunrise? The goal for Sunrise is a mark holder wants to be on a registered domain name during the Sunrise period of a TLD. It's pretty straightforward, it's pretty simple. I've got a mark and I want to be able to use that mark to register a name during the Sunrise period of a TLD.

The challenge with that, you must have an eligible mark registered in the Trademark Clearinghouse. And the registry needs to be able to verify that. So that's the two things that need to be able to happen during Sunrise, in the Clearinghouse, your mark has to be in the Clearinghouse and eligible. And a registry needs a way to verify that.

So that's the challenge. The requirements of absolution that we came up with anyway is that it needs to be simple, it should be decoupled, it

needs to be supportable, and it needs to be respectful of existing processes.

But there will be trade-offs between those requirements. You can't always meet every requirement. Some particular solutions might be really simple and decoupled but they might not respect existing processes. Another solution might respect existing processes but be really complicated.

There's always trade-offs to be made. We need to put forward the one that best meets the most of the requirements that we can. Who are the people involved? There're really four actors in our simplified model: the Trademark Clearinghouse, the registry operator, the registrar, and the mark holder.

There are other more complex models out there. There are resellers and so forth. Our model takes into account those other situations. But for the purpose of explaining how it works, we're just going to consider those four actors.

At a high level, what's the normal process that takes place? Normally a process that would take place regardless of the implementation is something along the lines of this. the Trademark Clearinghouse will collect and verify the information. It will store it in a database somewhere.

At some point the mark holder will request registration of a domain name in one of the TLDs. The registrar will send that request to the registry. The registry needs to confirm those mark details. Then the name is allocated over everything is okay and correct.

That's the normal process that we need to be going through during a Sunrise. Now of course it's a bit more complicated. There's probably a period where things are open for registrations. There might be auctions at some point if there are multiple allocations, etc.

But in the simplest view, that's all that needs to happen. So, at step five the registry needs to know something that the Trademark Clearinghouse knows. At step five in the process, this we're going to take me back, yes.

Step five in the process, the registry confirms mark details. At that point the registry needs to know some information that the Trademark Clearinghouse knows. We need to find a solution so that we can make it so that the registry can know what the Trademark Clearinghouse knows.

So how do we do that? Well, the simple solution is you just ask the Trademark Clearinghouse. The registry turns to the Trademark Clearinghouse and says, "Hey buddy. Do you have this mark in your database and is it eligible for Sunrise?"

That would be the simplest solution. The registry turns to the Trademark Clearinghouse and says that. But that doesn't meet our requirements. It's tightly coupled. If I turn to the Trademark Clearinghouse to ask the question and say, "Hey is this mark in your database?" and the Trademark Clearinghouse is busy, or is not there, or he's gone on holiday, or he's broken, he can't answer my question.

So we can't really have that situation. We need to look at a better way of doing it. How else can we do it? Well, that's where PKI comes in.

Most of you are probably going "Huh? What's PKI?" Let's hopefully try and explain it in a simple way.

PKI is complicated. But for our purposes the only thing we really need to care about is that it's a technical detail that the registry and the Trademark Clearinghouse can deal with. But for our purposes all we need to care about is that in PKI.

One entity can assert that a piece of information came from them and another entity can verify that that information did indeed come from the first entity. And that it hasn't been modified by a third party. So that's all we really need to know about PKI.

That's what PKI does for us. So it's technical and it's all to do with the wonderful world of mathematics. It's basically fundamental to how the security on the Internet operates today. It's technology that's been around for a long time.

It's the same technology that protects you when you log in to your Internet banking website. Or any time that you go to a HTTPS website, a secure website and URL. It's the same technology that takes care of all of that.

So how does it do that? Well, I'll try to explain very quickly how it does it. There are two actors involved and I call them the asserter and the verifier.

So the asserter generates what we call a public private key pair. There are two little bits of digital information. One is called the public key. One is called the private key. They're mathematically related to each other somehow, it doesn't really matter why.

The public key is, as the name suggests, public, distributed to the world. Everyone can know it. It's not that important but it's associated with the private key.

That private key is kept secret and it's protected. The Trademark Clearinghouse itself will hold on to the private key. So because of this wonderful mathematics that sits behind of this, the entity, the asserter, they can use the private key to what we call digitally sign some data.

They have some data using the private key and some algorithms and a bunch of mathematics they come up with, what we call a digital signature over that data. Then anyone, the verifier, anyone who wants to verify it, can use that public key, the one that can be known to the world, to verify the digital signature over that data.

If the signature verifies, then we can assert that the information that we have came from that entity. It's the information that entity intended us to have. Whoever holds that private key; this is the information that they wanted me to have.

And the digital signature can be used to verify that's what we wanted. If the data that I've been given as the verifier was generated using someone else's private key, or it was modified after it was signed, either mistakenly in transport or deliberately by somebody trying to pretend that they're somebody else or so forth, the signature will no longer verify.

So that digital signature that was generated using the private key by the asserter will no longer be able to be verified by the verifier. With this

assurance we can know that the data that we're looking at is the data that the entity that signed the data intended us to see.

All that anybody needs, anybody that wants to verify that data, all they need to do to be able to verify the data is to be able to have access to that public key.

So, in summary really quickly, private keys stay private and are used to sign data. Public key is made public and is used to verify data. How does this apply to the Trademark Clearinghouse and Sunrise? Well, hopefully it's pretty straightforward and you guys see how it's all coming together.

The Trademark Clearinghouse will generate a public private key pair. The Trademark Clearinghouse will distribute the public key to all the registries. All the registries will have the public key.

Public key is a tiny little bit of information, probably less than the size of a picture taken from your camera. Each registry will get a copy of that key. In fact, the whole world can get a copy of the key. It doesn't matter. It's public information by its very definition.

The mark holder will register their mark in the Clearinghouse. The Trademark Clearinghouse will give the mark holder back to them what we're calling in our model the signed mark data, the SMD if you've read our documents.

That's simply a file that contains a subset of the record in the Clearinghouse that's been digitally signed by the Trademark Clearinghouse's private key.

So a file, subset of the information out of the information that the mark holder put into the database, signed with the Trademark Clearinghouse's private key. We call that the SMD, the signed mark data.

Then when the mark holder wants to register a name during any Sunrise in any of the TLDs, all they'll do is they'll provide that SMD to the registrar that they are using. The registrar will pass on that SMD to the registry as part of the normal domain create command and the registry.

Using the public key that we have that the Trademark Clearinghouse gave us, we can verify that the information that we got from the registrar, which the registrar got from the mark holder, is the information that's actually in the Trademark Clearinghouse. And it has been verified and approved by the Trademark Clearinghouse.

And because of that PKI model, if anybody tries to modify that information, either mistakenly or deliberately, the signatures will no longer validate.

The registry will know that this data is either data that didn't come from the Clearinghouse or data that has been modified by somebody, or whatever it is that happened. So if we validate the signature and the signature verifies, the registry will allow the registration to occur.

Now whether that now goes in the queue for a Sunrise process for some auctions or something, or it happens in real time, that's up to the registries to figure out. So the only information that was required to be provided by the Trademark Clearinghouse to the registry was the public key.

There was no sharing of data between the Trademark Clearinghouse and the registry. The Trademark Clearinghouse gave the registries just the public key. All the other information that the registry receives, they get it from the mark holder directly.

The mark holder makes a choice to provide the information to the registry to participate in the Sunrise. If they choose to do that, then the registry can use the Trademark Clearinghouse public key to verify that that data has been verified and accepted by the Trademark Clearinghouse.

It's all pretty simple and straightforward. It's not too complicated at all. This meets our goals, our requirements. It's simple, it's decoupled. The Trademark Clearinghouse can disappear. And for as long as we want afterwards we can still verify these signatures using the public key. It's supportable. These signed bits of information are able to be read by registrars and registries.

And it's respectful of the existing processes. It meets our goals. It also has some other benefits. Other benefits of the model, this includes things like registrars are able to detect and fix issues.

So because this signed marked data information is human readable, if you accidentally try to register one of your other marks in the Sunrise utilizing the signed marked data from a different one of your trademarks?

The registrar will be able to detect that issue and let you know and tell you that this is the wrong information. You gave me the SMD for the fu trademark and you're trying to register the bar name. Mark holders are

in control of their data so there's no arbitrary sharing of data between the Clearinghouse and registries and so forth.

That's all within the hands of the mark holders. Registries can request data from mark holders so there's no independent validation of that data required. One of the goals of the Trademark Clearinghouse was to try.

And in Sunrise, the past mark holders participating in Sunrise would have to submit their trademark data to various registries. Then they would pay a verification fee to each of those registries. So that those registries could go and verify all that trademark data using whatever way that they did that.

Most of them outsourced it to various other entities. The whole point of the Trademark Clearinghouse is that that validation gets done once by the Trademark Clearinghouse. And then each of the registries can rely on the fact that the validation has been done by the Trademark Clearinghouse. And not have to charge those expensive verification fees over and over and over.

By sending this signed mark data information to the registries the registries know that the information is trusted because it's signed by the Trademark Clearinghouse. They don't have to get independent verification of that data.

So if a registry has a policy that says you have to have a trademark in a certain class of goods to participate in my Sunrise, they don't have to ask you for that information and go and get it independently verified. It's already been verified by the Trademark Clearinghouse.

That signature tells us that the data has been modified and correct. It uses proven technology. As I said, this is the technology that makes the Internet work today. This is how all security on the Internet works. There's nothing new here.

We're not inventing anything and there's less chatter. There's less back and forth. There's less information being passed between Trademark Clearinghouse and registries.

In fact, there's only one piece of information and that's the public key. That's it. That's the model. That's the registry registrar community model that we're suggesting. Yeah, I'm happy to take questions if there are any.

KAREN LENTZ: Thank you Chris. Would you like to come up?

MARC TRACHTENBERG: Hi. I'm Marc Trachtenberg with Winston and Strawn. What data if any is in the key itself?

CHRIS WRIGHT: In the public and private keys?

MARC TRACHTENBERG: Exactly.

CHRIS WRIGHT: Yep, there's no data in the public and private keys. The public and private keys are just a binary string of numbers.

MARC TRACHTENBERG: So I guess I'm just not really clear on how the data can be verified just by the key itself. How do you connect the public key to a particular mark?

CHRIS WRIGHT: So the data, the trademark data, the information that this trademark, this is the trademark and these are the DNS labels that it corresponds to and so forth. That's in the signed math data. That's in the information that's been signed by the keys. It's not the keys itself. So there are only two keys. There's a private key held by the Trademark Clearinghouse. There's only one of them that it has.

MARC TRACHTENBERG: So I'm the trademark holder. I submit my information to the Trademark Clearinghouse.

CHRIS WRIGHT: Yes.

MARC TRACHTENBERG: They create this key.

CHRIS WRIGHT: No, no. the key is an independent part. So you're referring to the signed mark data, the information that gets signed. Yep. There's one private key and one public key.

So the private key is just a string of ones and zeros that the Trademark Clearinghouse holds and he keeps it secure and protected. You submit your data to the Trademark Clearinghouse. The Trademark Clearinghouse generates a signature over your data using the private key.

Then you get back your data with the signature on the bottom of it. The key stays with the Clearinghouse, the private key, and the public key the registries have.

When you send your data to the registry via the registrar, the registry can verify that signature on the bottom of the data using the public key. That signature corresponds to the information that's in that piece of signed marked data.

MARC TRACHTENBERG: Okay.

CHRIS WRIGHT: Does that make sense?

KAREN LENTZ: So we have a slide that might help visualize this a little more. And I'm going to have to go back because I messed up when integrating the two

slide sets. So a Sunrise code, which is what was in the Draft Implementation Model that ICANN did looked something like that.

When you're going to the registrar, you're providing that code that demonstrates that you're eligible to registrar on this Sunrise. A signed marked data file would look something like, and Chris you can tell us if this is right, would look something like that.

CHRIS WRIGHT: In one method, yes.

[background conversation]

MARC TRACHTENBERG: So the Clearinghouse provides back a file that it has signed with a key?

CHRIS WRIGHT: Yes.

MARC TRACHTENBERG: That's the part, that's the disconnect I was having. So you have a data file they're giving to you and you have to submit that data file with your Sunrise application.

CHRIS WRIGHT: Correct.

MARC TRACHTENBERG: Okay, that was the part that I was missing. Okay, thank you.

KAREN LENTZ: Yes. And then in this kind of for purposes of fitting on a slide, it's all compressed together. But you could organize this so that it made more sense with the tags. There's jurisdiction. There's mark.

There're various fields that are in here. So the questions that we wanted to pose in terms of the Sunrise, we looked at the PKI model that was proposed as Chris described. We think it technically can work.

We think the codes can work and the PKI system can work. We wanted to have a little bit more of a discussion around a couple of issues. And I see there's a line so we'll get to you.

One is usability. Is there a difference in the provision of a code versus a file? The other is security type issues. Is there a difference if a code or a file is lost, stolen, or needs to be replaced for whatever reason? Then as well the SMD can be configured so that it has a minimal set of information or a maximum set of information as far as what's in the Clearinghouse records.

These are the questions we wanted to pose to all of you as potential Clearinghouse users. So we'll go with James.

JAMES BLADEL: Thanks Karen, James Bladel with a comment and question. I was there in Brussels when this started to take shape. So I know that the registries and new applicants like this community model.

Speaking on behalf of a registrar and with some assumptions that this is operationally as reliable as any other registry system that we've put out there. And that the DNS is relying on today. But understanding that that's one caveat that this can scale and it's responsive, etc.

I know that we like it because we have to explain it to our customers and we have to support it and answer their questions and help them find and reset their lost keys, etc.

A few of the folks that I know would be consumers of the Trademark Clearinghouse in terms of marks holders also prefer this model because it doesn't expose their database of strings and jurisdictions, etc.

it seems from their perspective to be more secure. So my question is who doesn't like this? It seems like we've got one of those lightning in a bottle moments in ICANN where everybody likes something.

And my question is who doesn't? And Christine's behind me. Maybe, Christina are you going to say you don't?

CHRISTINA ROSETTE: I'm going to say not everybody's made up their minds yet.

JAMES BLADEL: Okay, that's fair. But it seems like we're really getting kind of close to one of those rare occasions so I guess, what are the criticisms? Help me understand please. Thank you. I'll sit down.

KAREN LENTZ: Christina?

CHRISTINA ROSETTE: Yeah, Christina Rosette from Covington. I just have a clarifying question. Just so that I'm clear under the original draft model, a trademark owner would receive a different Sunrise code for every single trademark registration that they deposited with the Clearinghouse.

KAREN LENTZ: The way it is currently is per mark.

CHRISTINA ROSETTE: Right, so if I have mark AB in Japan, I get one code. Mark AB, same mark in Canada, I get a different code. Under this model will the trademark owner get a different signed data file for every trademark registration that they deposit? In other words, will they still have multiple pieces of information that they need to track? It'll just look differently.

CHRIS WRIGHT: You can do it whichever way. You could make one giant signed mark data file that had several marks in it or you could have one piece of signed mark data for each mark. It can support whatever you would like.

CHRISTINA ROSETTE: And when you say whatever I like, does that mean on the per user basis or is that something? It sounds as if it's something that would have to be decided pretty much kind of top down so that everybody is doing the same thing.

CHRIS WRIGHT: Yeah, it would be easier if we did it that way. Yes.

CHRISTINA ROSETTE: Okay, thank you.

KAREN LENTZ: Next, Jonathan?

JONATHAN ROBINSON: Hi, it's Jonathan Robinson. I've participated in this process in the technical process on behalf of affiliates being in the Brussels meeting and various emails. I just wanted to make sure. Mine was really a clarifying statement, if you like, connected to Christina's point.

I just would make really sure that in the prior ICANN model if you like that the data was. I have an image of data bundled up in an encrypted bundle that's transmitted to be used for a single purpose use.

In the PKI model that Chris has described, the data is transmitted in plain text, signed plain text, correct? And it can be used for then multiple purposes. I think one of the attractions is that data is available for various different purposes.

So I think you explained that it can be used for class of goods, country, different variants, different variants of Sunrise could then be applied to that same data and used. So there's different models of Sunrise can be used.

CHRIS WRIGHT: We need to be careful. We seem to be confusing a little bit here the claims process with the Sunrise process. In neither this suggested model or the ICANN model, well actually the other way around. In the ICANN model suggested for Sunrise there was no mark data being transmitted to the registries at all.

There was only a single code that somehow the registry could verify using a similar sort of hashing technique. But that code corresponded to a particular DNS label.

The difference here is that in that model it was difficult to see how the registry would be able to get access to the other information about the mark to perform.

Perhaps they're doing hierarchical allocations. For example, I know of a new TDL applicant that intends to in their Sunrise have hierarchal based allocation rules where it's a GOTLD.

They intend to say that trademarks registered to organizations in that country get preference in their Sunrise over trademarks registered to countries in their region that have preference to trademarks registered to all over the world.

So that's their policy that they're looking at applying where without having that information they're not able to make such hierarchal based allocations for example.

KAREN LENTZ: Okay, thanks. We'll go over here to Rubens.

RUBENSS KUHL: Rubens, NIC.br. Just to comment on whether the information is encrypted or not, the register to register communication is usually encrypted by TLS and the user to register communication can be encrypted as well. It's just a matter of the register employing its HTTPS on their site. The flow information can be encrypted as well.

CHRIS WRIGHT: Yes.

RUBENSS KUHL: But I would like to make another comment is that we should offer peer user options of having or not class of goods and countries listed at the SNB.

Because if we try to get a one size fits all model we always have someone, "Hey, I have a privacy concern. I don't want to disclose which country my mark is registered. I don't want to disclose that my mark has that class of goods."

And last it has to require that is an eligibility requirement so we should probably offer [opinions] so you cannot need to go to dive into those issues. Let the users who are the data owners tell us what they want.

CHRIS WRIGHT: So a couple of things. I understand what you're saying. The reason we discounted that from our solution was that starts to get really complicated now because now you're saying, "Hey mark holders, you need to go and generate different types of SMDs for all the different TLDs that you're going to go and register in."

So now we're trying to remove the headache for mark holders here. We're trying to make it so you download one SMD file. You use that across every TLD that you intent to interact in.

That's why we discounted that. But if mark holders are going to tell us that there really is a big privacy concern with telling a registry that my mark is registered in this country, then so be it.

Then we can go down the model of allowing a whole bunch of different SMDs to be generated. But at least intuitively we thought that organizations with lots of trademarks it would be a bigger headache for them to try and deal with different SMD codes for all different registries than to just accept that I'm telling the registry that my mark is in a particular country. I'm not really sure that we get what the issue is with that.

KAREN LENTZ: Thanks everybody who commented. I want to speak a little bit to the person who asked who is against this, or what would be the possible criticisms to it. I'm not meaning to present these as criticisms but these are just some of the questions that were raised in the discussions.

Because the SMD file does have plain text data in it, it means that's being transmitted around. So if it gets into someone's hands that wasn't supposed to have it, is that a concern.

Then one of the points that I think was made in the proposal was that this may not be a big issue in Sunrise. Because the rights holder is only giving data their registry because they want to register a domain name there's nothing extra in there necessarily.

But part of what we did want to get comment on was first that issue of having the data in plain text. Secondly, what's the right set of information to have in there?

Are there different views on what the appropriate content would be for a signed mark data file that's used for Sunrise? So we'll go to the next speaker.

CELIA LERMAN: My name is Celia Lerman. I'm from PIR. My question has to do with the public part of the key. I'm looking at this like an SSL certificate. So you've got your public and your private, if somebody else gets hold of the public key.

KAREN LENTZ: Sorry, can you speak up a little? I think people are not hearing you.

CELIA LERMAN: Oh, sorry. I'm too short.

KAREN LENTZ: You can, it moves.

CELIA LERMAN: Okay, thank you. If someone else gets hold of the public key. That was longer than my question I think. If somebody else gets hold of the public key, how do you know that's really the trademark holder that's

submitting it? And are these updated periodically and if they are then the trademark holder has to continue to get the new updates.

CHRIS WRIGHT: Yep, so it's similar to the question that we had at the start. There's only one public key. There's not a public key per mark. There's not even a public key per registry. There's only one public key and it's the Trademark Clearinghouse's public key.

The public key is what you use to verify the signed mark data. So that's that XML example that was up there before. If I get your question right, what I think you actually mean is what happens if somebody gets hold of my SMD.

Well it means they can go and register a name using that SMD. But in the ICANN model that was presented, what happens if somebody gets hold of your code? They can go and register a name too.

So either way a mark holder has to protect something. Either they have to protect a 16 digit code or a 32 digit code or whatever it is. Or they have to protect an XML file.

Now the reason that we prefer the SMD is because if this is really an issue, then what we can do is we can add more fields to the SMD, such as the name, address, and telephone number of the person that's allowed to use it.

Then if somebody else is magically able to get ahold of your file somehow, all they can really do is register you a free domain. So that's

why we prefer this model as well because it actually protects against this sort of problem.

As long as mark holders are now comfortable with putting their company name and address in there for example, it's always a trade-off. The more information you're prepared to put in there, the more secure off it will be. But of course that's more information that's in there that maybe people don't want. So it's a trade-off.

CELIA LERMAN: Can I have a second question? So, will each registry then decide what needs to be in that mark or is there going to be a general decision about what's in these SMDs?

CHRIS WRIGHT: So that's again similar to the question from before. We can go both ways. We can make it so that each registry can say that if you want to register in my Sunrise, your SMD has to contain these fields. And another registry could have a different set of requirements.

Now the trade-off there is that if you go down that method, as a mark holder I have to deal with perhaps five, six, seven, who knows, a thousand different versions of my SMDs and keep track of which one belongs to which registry.

If you go down the other way, you can just have the SMD with the agreed set of data that everyone would put in there. We'd have to work out how we figure out what that agreed set of data is. But we figure out

what that agreed set of data is and then you only need one and you can use with every registry.

So there's trade-off both ways. You guys have to deal with this more than we do as registries. So we're really sort of pushing this back on you. You guys can decide if you're really worried about all the amount of data exposure.

And you're happy to deal with all these keys, all these different keys, then so be it. go for it. It's no skin off our nose. At the end of the day, we get a bit of data that we verify. As long as it has the bits that we need in our TLD, we're happy. So whatever you guys want.

CELIA LERMAN: Well I think to simplify it for the mark holders, they would want one set. Which I think would be, since there's potentially 1,400 TLDs coming out, they're going to want one set.

CHRIS WRIGHT: I think so too.

CELIA LERMAN: But I think it would be up to the registries to get together perhaps and figure out what needs to be in that SMD so that everybody knows it's pretty secure.

CHRIS WRIGHT: So the registries could propose we think this data needs to be in there. You guys could come back and say we don't want that field in there. We

could talk about what the ramifications of that would or wouldn't be. I think we, where's Geoff? I think we would be able to do something like that, right?

GEOFF BICKERS: Yeah.

KAREN LENTZ: Jonathan?

JONATHAN ROBINSON: Yeah, thank you. It's Jonathan again. I'm just going to make a very brief comment. Chris made reference to the fact that the SMD data could contain various options.

For example one option one might think of that will need to be decided is whether the trademark is registered to the mark owner or someone else's data within the Clearinghouse. That's essentially a policy.

I'm not sure if that policy is finalized at this point because that might apply to all domain names registered on the back of that mark. But those are the kinds of variance of data that could be put into.

Is there a registrant in addition to the mark owner, for example? That could be contained in the SMD data. So without wanting to go into whether that's the right policy or not, it's an illustration of what an option could be within that file.

KAREN LENTZ: Okay, thank you. J. Scott?

J. SCOTT EVANS: Yeah this is J. Scott Evans from Yahoo! And I'm not a technical person but I'm a trademark attorney. I have the paralegals I have to explain how to send all this information in.

So my question is, goes to Christina as in this question about what fields go in. On day one we have five registries that are going live. Do we have to put every trademark in our portfolio and every jurisdiction in on day one? And continue to do that in order to participate?

Are we going to get warning time so they're like the New Zealand registry is going live? They only want to do New Zealand stuff. So I have time to decide whether I want to put my New Zealand marks in because that is a huge issue for budgets and whether we're going to participate.

It has to do with what fields are going to be in because the New Zealand registry may, for that round. Those groups may have different fields they want to put in to their format. So that's the question I have.

Is it all in now and every time I register a mark I just have to just also put in the Trademark Clearinghouse? Or is there going to be a step that there's going to be some sort of notice that xyz registry is going to go live. You can decide whether you're going to participate.

If so, you have to go in to the Clearinghouse at this time and you have to provide this information.

CHRIS WRIGHT: I think there are a lot of different facets to that question. On the first part, when do I put my marks in the Clearinghouse? That's actually completely up to you when you put them in there.

So when you put a mark in the Clearinghouse, if you want to participate in the Sunrise, your mark needs to be in the Clearinghouse. As for notification when a Sunrise is going to commence, well I think that's going to be left up to each registry to do whatever it is that they feel fit.

That's a competitive issue amongst the registry. If they don't give you enough notice, you're not going to register, and they're not going to get registrations.

Another part of your question was kind of leaning towards I already have my mark in the Clearinghouse and I only put certain data in the SMD. Now a new registry is going live and they need additional data that I didn't put in the SMD.

So that's kind of different to I didn't have my mark in there yet. And presumably to do that, if we did go down this model of having different SMDs with different amounts of data in them which my personal opinion that's probably a little crazy because it's going to get to be a headache for all of you.

But if we did do it on that model, presumably you could just go back to the Trademark Clearinghouse website, seek a few additional boxes to say generate me a new one that now included these data and away you go. I don't think that would be too difficult.

What was the last part? I extracted a third part from that somehow. But that's just another reason why having different bits of data in there for

different registries would probably end up making it harder for you guys.

KAREN LENTZ: Okay, thanks. So we want to kind of wrap up this and get into claims as well in this session. So we'll go to Tom and Claudio and then we'll go on.

TOM BARRETT: Tom Barrett from Zurka. I think one point that maybe you haven't talked about yet is the fact that during Sunrise not only are we to submit an SMD object, but also the normal whois contacts associated with the Sunrise registration.

So the question is you have an object, you have a registrar admin tack and billing. And in addition to verifying the SMD data itself with a Clearinghouse, is there also the need to verify the whois data somehow matches the SMD data? Or are we allowing those to be decoupled and not need to correlate?

CHRIS WRIGHT: That's a question that we had. Again, we don't know the answer. It's another one of those policy questions. Realistically what it boils down to is if I put trademark x into the Trademark Clearinghouse with company details, Company ABC.

If somebody then tries to register a domain using trademark x and its SMD and in the registrant field they put Company XYZed, should registries not care. Should registries be rejecting that? Should registries, I don't know. What should registries be doing?

What's the expectation of what registries should be doing in that situation? For most of us, we'll probably just let it go through unless you tell us we shouldn't be doing that.

If you do tell us we shouldn't be doing that, then what we would be saying is then put the company name in the SMD. Because then we'll just extract it from there and use it. But it's again; you guys need to tell us what you want.

KAREN LENTZ: Claudio?

CLAUDIO DE LUCA: Thanks Karen. Thanks Chris, again, for the presentation. Karen I wanted to thank you. I know this has been kind of like a pet project of yours and there's been a lot of progress made. Clearly a lot of issues involved, we think the Clearinghouse is critical to the entire New gTLD Program. So we think it's really important.

I wanted to confirm if there are plans to post this for public comment or the ICANN implementation model for public comment? I know there are a couple narrow issues that are posted now about proof of use and matching. The IPC is going to be submitting comments on those.

But it's difficult to look at those situations in isolation from the overall model that ICANN is proposing. Also what this alternative proposal is about.

So we're going to be, in our comments we propose that the comment period be extended for at least 30 days so we can have really a proper

consultation. These are great sessions but the affirmation of commitments really requires, we believe, really a robust consultation.

So we're encouraging the extension of the comment period. I just wanted to confirm if that's something that you guys have planned or how you might move forward on this.

KAREN LENTZ: Thank you Claudio. So two things, on the comment period is really on the two memos that you mentioned on the proof of use and matching rules which are pretty much independent of what you do on this end with the registries.

In terms of having a public comment period, I think that's potentially one of the outcomes of this meeting. Although I think we're a little beyond here's a model, here's another model, pick one.

We're trying to get into what are the concerns that people have raised and what is the sense of what the right way to go here is. Then we can get more into building it.

It's certainly possible. What a comment document would look like, it might try to distill some of these issues and frame them but thanks for the suggestion. All right, Christina, real quick.

CHRISTINA ROSETTE: I just have a request. I think speaking for the IPC it would be super helpful for ICANN to let us know what you need to know from us. I mean a very specific list of questions, as detailed and specific as you want.

The more you ask us to tell you, the more we're going to tell you. We're happy to give you the information that you want once we've decided it and developed it. But I want to make sure we're you what you need and not, not enough.

KAREN LENTZ: Great, thank you. All right, so to sum up Sunrise I didn't hear anyone really violently objecting to either approach. There are some questions about how to configure it so it's the optimal solution for the registry and for the trademark holders, user.

So we'll continue to consider that. Going to the trademark claims, I'm going to go back here just a little to the beginning because I want to have a little high level discussion first.

The trademark claims service, just to review it real quickly, what it's intended to do is for a 60 day period when a registry is in its startup phase.

I go to register a domain name. There's a check down against the Clearinghouse records to determine if there's a trademark record that matches the domain name I'm trying to register. If there is, there's a notice generated that says, "Hey, you should know there are these marks in the Clearinghouse. Do you wish to proceed?"

I decide whether I want to proceed or not. If I do, I must acknowledge that I've seen that notice. Then that gets transmitted back. And the Clearinghouse notifies the relevant rights holders that domain name has been registered.

So this is a difficult process to build, maybe not on the level of whois. But it's hard so we've struggled quite a bit with how to build this in a way that makes sense and is useful for everybody.

The draft model that was produced in April had this implemented by distributing the Clearinghouse data to the individual registries in an encrypted form really to provide some accountability as to the use of the data.

In response to some concerns that had been expressed regarding what could somebody do with the aggregation of all the records that are in the Clearinghouse? It provided that the data would be sent in encrypted form.

The necessary information would be passed through the registrar displayed to the potential registrar and then passed through so that the notice could be provided. Now, there's an alternative claims model as well, which Chris or Geoff or someone can explain better than I can. Geoff, would you like to give a few comments on this?

GEOFF BICKERS: Thanks. If we can do those slides maybe we can walk through it because it's easier to kind of see the picture. I actually gave Kurt, Kurt gave them hopefully.

[background conversation]

GEOFF BICKERS: Bring up Geoff's slides please. How's that, good? Good, then we can see it? If we go to, actually go to the next one. I think it's on the, start on the next one. There we go. So I thought we'd kind of walk through the claims process.

There are a couple differences between the centralized model. Let me go back a step. The proposed model from ICANN we call the decentralized model because as Karen said, it really pushes all the data out to the registries.

Then it really relies on the registries to do the bulk of the work, meaning the matching and the sending out the notices. The model, the alternate model or the registrar and registry community model really is what we dub a centralized model where the registries will be doing some of the checking.

But the bulk of it is actually going to be done at the Trademark Clearinghouse level. So under each model, some of the slides actually apply no matter which model. I'll walk through the process and I'll let you know where the two models actually differ in terms of how it actually gets implemented.

So the very first thing, whether it's the centralized or decentralized model, is that the Trademark Clearinghouse will, and I use the term loosely, publish the list of strings.

In the decentralized model, the ICANN proposed model, it's actually the entire Trademark Clearinghouse data. When I say published I don't mean published for the world. I know that's kind of a loaded term. Published means just distribute to the parties that need it.

In our model, the centralized model, even in that approach the registries would still get a list of all the strings in the Clearinghouse. Not the trademarks which they're based on, not the owner of the trademark, not the jurisdiction, or any of the other data that is in the Clearinghouse. Really it's just a list of these, all the strings in the Clearinghouse.

See if I can go and do this, there we go. So we've broken this down to the registrant, registrar, registry, and the Trademark Clearinghouse. So the registrant, we finished the Sunrise period and now we're into general availability.

The registrant goes to register a domain name, goes to their registrar just as they do in the normal registration process. The very first thing that's done, which is done in any registration process is the registrar says, "Okay is the name that the registrant wants, is that name available?"

If the name is not available, which we didn't include on this chart, then obviously the registration process goes no further because someone else has it or it's on a reserve list.

But if the domain name is available then we go back to see if the domain name subject to a claim? If it's subject to a claim, the process would then say that the information would be returned.

But in this case, the registrar would ask the registry, "Does the name match a claim?" if the answer is no, then it just goes through the registration process as any name would because there's nothing in the Clearinghouse that matches it.

Let me actually go on to the next one here. Okay, so let's assume that there is a claim. This is the slide were there are some things that are different between the ICANN proposed model and the alternate community model.

In this slide, if you look here, what I'll go over first is the ICANN proposed model. Then I'll talk about this one that's up on the screen. So in the ICANN proposed model, if there is a claim, since the Trademark Clearinghouse has sent all of the information through the registry, to each registry, that box that's up on the right hand side that says, "return claims notice".

Just ignore some of the abbreviations because this was done for a different presentation. That box that's now under TMCH in the ICANN proposed model would be under the registry so there never would be a need in the ICANN proposed model to go to the Clearinghouse to get any information.

Ultimately in any claims process, whether it's our model or the ICANN Model, information would be returned to the registrant in the form of a claims notice that would say, "Okay, you've applied for xyz dot web and that matches a mark or one or more marks in the Clearinghouse.

Here's all the information about the trademark that you'll need to make an informed decision to decide whether you want to proceed or not." If you do want to proceed and what we have here is the accept the notice. The registrants would say, "Yes, I accept."

Then that acceptance gets transmitted to the registrar, in both models, and then that acceptance goes from the registrar to the registry and the name is registered.

So the core fundamental difference between the two models is who does that function that's in the top right hand corner, which is returning the claims information ultimately to the registrant. Is that done by the registry or is that done by the Clearinghouse in one centralized place?

Go on to the next one. So we talked about this. Essentially this next slide is what information is returned when someone wants to register the domain in either model, or both models.

The registry will capture information such as yes the registrant was displayed the notice and the timestamp of when they accepted. So certain information the registry will capture in order to, I'll use the word prove, I know some technical guys will get on me for that.

But basically to prove that the registrant was displayed the notice and that they proceeded anyway. In the model, so there are two models that we talked about.

The whole question is, is the notice being provided directly by the Clearinghouse or is the notice being provided by the registry? In the ICANN model the entire trademark claims database in an encrypted form would be sent to each registry.

And for a number of reasons that have been argued, regardless the type of encryption that's used, I'm not sure how to say this without scaring too many people.

But regardless the type of encryption that's used, essentially it'll be fairly easy for any of the registries to decrypt that information and to get copies of that database.

So the real question here is how comfortable is this room and the IP community with the notion of registries being handed all of the information.

In the centralized model that we have proposed, there's one source of that information and the only information that the registry gets is actually really no information.

Essentially in our model, and I'm trying to put this as high level as possible here. In our proposed model, the registrant goes to register a name, sends its information to the registrar.

The registrar checks to see if the name is available with the registry. The registry says yes. Then it asks the registry, "Is there a claim that matches it?" the registry will check; now this is the part that the registry, the registry has the list of strings.

So the registry can check its own database to see if there's a claim. If there's not a claim, it's pretty easy. The registry just goes forward and registers the name.

If there is a claim, then in our model the registry would say, "Please Mr. Trademark Clearinghouse, or Miss Trademark Clearinghouse, send me the information for the claim so that I can send that on to, or the registrar can send that on to the registrant to give that information."

So that's essentially the two models compared back and forth. There are a lot of pros and cons to each approach. For example in the decentralized model, in the ICANN model, it doesn't matter if the Trademark Clearinghouse is up and running because all of the data has now gone to the registries.

And really if there is a "failure" in the process, it's really one or two registries that may fail. [That] actually correspond to a claim will not be able to move forward. So let me give you an example and say that in a different way. In the centralized model, you go to a registrar [to complete] a registration.

However, in the proposed model we have come up with a number of ways to mitigate that. That mitigation is important because we believe based on past experience, now there's only been one trademark claims process ever launched. That was in .Biz in 2000, 2001.

Now a lot has changed since then but during that time period we, and this is public information that we filed in a proof of concept report, about ten to fifteen percent of the domain registrations that were registered in the first 60 days actually matched a claim.

Now you could say that's only one TLD so that figure might be low. But we would estimate on a very conservative that at the most you're talking about in the first 60 days of registration you're talking about maybe 15%-20% matching a claim.

So in our model in which the registry actually has a list of the strings will be able to tell whether your registration or registrant's attempted

registration matches a claim. If it doesn't match a claim, which we believe will be at least 80% of the time, and then it's great.

It doesn't stop any registrations from going through, those 80%. They'll go through. They never matched a claim, it doesn't matter. We're only talking about the percentage of domain name registrations that actually match a claim. Those are the ones that will be impacted.

For that, in our model, we have proposed just reasonable measure that any entity that's running a database like this, a live query system, should take into account things like having some redundancy, living up to some SLAs, having a reasonable business continuity plan, having data escrow.

I almost said it like I've been listening to the Australians. I almost said data escrow. I've been with Chris a lot.

So reasonable precautions so hopefully that wouldn't happen. And as Chris said in the Sunrise slide there are trade-offs. There are trade-offs to either approach.

What we're really reaching out the community because what we were told when we were developing this that number one concern from trademark owners, or one of the top concerns from trademark owners was that they told us that they did not want all of the data to go in bulk to anyone outside the Clearinghouse.

That was a very big concern. If you go with the centralized approach, which is the one that we are proposing, that concern becomes mitigated because it's really just that one entity that has all of the data.

If you go with the decentralized approach, then you are submitting that data to a number of different entities, you're asking those entities, all of the registries, to write their code on top of it to implement all of the distributing notice.

And you're basically relying on those 1,400 or depending on the number of backend providers, a certain number of registries to actually implement it. However, the trade-off is if you go with the centralized approach.

As I said, you're relying on one entity. If that entity is down, then for a certain percentage of the registrations they're not going to go through. Again, it's trade-offs. The question we're asking to the community is what are the most important values to you?

Are you willing to live with the trade-offs? Because if you say, if the IP owners say, "You know what, we're comfortable if the registries get all of the data and we're comfortable with all of that."

Then maybe the decentralized approach is the way to go because it does mitigate the risk of having one entity go down and stopping registration.

If you tell us, on the other hand, "No, I really think that the most important thing is not to distribute the data." Then the centralized approach is the way to go.

I know that's a lot to swallow. It's what we've been grappling with for a number of months. In this kind of scenario I don't think there's a right or wrong answer. Either way we can work with. I will tell you from the registry standpoint, kind of selfishly speaking.

As weird as this sounds you'd think entities would want as much data as they can get because we love having data. In this case, each of the registries has supported the centralized approach. We actually don't want the data.

Because if we get the data we know there's going to be expectations on us from the community that we protect the data. We safeguard it. We make sure there's no data mining.

It's much harder to do that with 1,400 players than it is with one. Again, there's trade-offs. I tried to summarize. I know it took a long time.

KAREN LENTZ: Thanks, Geoff . We've got a question here. I wonder if I could ask a question first. Just to clarify. In the model you described which we think is a nice idea was to sort of separate the yes/no is there a claim from the whole data. So there's this list of strings. Is that accessible just to registries or is that accessible generally?

GEOFF BICKERS: The way we've proposed it, and again it could be obviously subject to input. We actually proposed that that list be available to registries, but also be available to registrars.

Because one thing we've noticed is that registrars who are in the registration process and dealing actually with the end users. They may want to know even before they send it to the registry whether it matches a claim. There are certainly a lot of valid reasons for them to know that.

But of course now again you're talking about further distribution of data which may or may not be considered to be something that trademark owners want to safeguard. In our model we propose yes that it also be distributed to registrars but again that could be limited to registries. That's really, as Chris said, we could do it either way.

KAREN LENTZ: Okay, thank you.

STEVE LEVY: Hi, Steve Levy with FairWinds. In the decentralized system, how long does it take to propagate changes from the Clearinghouse? And is there a great concern for the registries having different data at different points of time?

GEOFF BICKERS: So I can address the last part first. Then I might look to Chris. In the decentralized model, one of the things that's going to have to occur continually, and it might be once a day is that each registry will have to download the list of additional marks that have been added to the Clearinghouse, all that information.

In a decentralized model that adds additional risk, because the file could be corrupted, you could have, even through no fault of the file being corrupted, each registry there could be a bug.

You never know about. So it is possible that each registry has different types of data and Chris could probably explain it better.

STEVE LEVY: It seems like thin and thick whois.

CHRIS WRIGHT: Kind of, yeah. In any distributed system where we're distributing data to a whole bunch of different entities, there's always the problem of entities being out of synch. The ICANN model as it currently stands; we have suggestions on how to improve that to try to mitigate those scenarios.

Not eliminate because it's going to be almost impossible to eliminate them. But mitigate those scenarios. Look, we think we will be able to make it so in 99% of the cases it won't be a problem. But in a decentralized model we always run that risk.

GEOFF BICKERS: One of the things you could do to mitigate the risk, for example, is when you're launching a registry you could say I'm only accepting trademark claims that have been entered into the Clearinghouse by x date. Maybe that's a week or two before you actually launch.

What it means is that any claims that come into the Clearinghouse afterwards, we're not going to generate a claims notice. But at least if that information goes out to the IP community in enough time for them to digest they'll know.

"Okay if I want to have my claim go out to everyone who wants to register .web, we know .web is launching in May of 2014 and the date says that I know I have to get all my claims in by April of 2014."

Yes you're going to anger some people that are late in getting their claims in but as long as they have enough notice, that's as much as we can do.

STEVE LEVY: Thank you.

CLAUDIO DE LUCA: Thanks Karen. Geoff thanks a lot. Claudio De Luca, member of the IPC. Geoff thanks for the work that NeuStar has put into this. It's clear you have heard some of the concerns from the trademark community. And there does seem to be some good flexibility aspects of this alternative proposal.

Unfortunately I'm not in a position to answer some of those specific questions that you raised. So Karen, again, I just wanted to please if you could put this out for public comment.

I think the ICANN model has been on the new gTLD microsite since April. And it hasn't been posted for public review and comment. We need to come to closure on this clearly.

I think what the community is putting forward needs to be properly considered. And the ICANN model needs to be properly considered in the aggregate.

Please if it could be posted for public comment I think it would be a great way of going forward. Thank you.

KAREN LENTZ: Thank you. One of the questions, and then we'll get to Wendy, that I did want to make sure we posed in here was something that Geoff eluded to and that was the relative value of the encryption.

As Geoff also said, we did also hear the strong concerns from the trademark IP stakeholders that they had concerns about the aggregation of the data in particular.

And what intelligence you might be able to gleam from that that you wouldn't be able to do some other way. Recognizing that as a concern we tried to put in some motivating elements.

So really with the encryption there's no way that a registry is going to accidentally access the data. It would need to be done deliberately. Then we anticipated that there would be some contractual terms around this as well as to what constitutes acceptable use.

Encryption is something that's been commented on quite a bit. It does add some cost and complexity throughout the system. So one of the questions we wanted to post was, are there other ways that the same types of objectives could be met? Could we do away with the encryption and what would be the impact of that?

GEOFF BICKERS: Sure, thanks, and I think it's almost, and Chris can jump in here too, it's almost it doesn't really matter what type of encryption you actually put in or how strong it is.

I'm not even necessarily talking about an unscrupulous registry that tries to get the data. Essentially trying to make it as high level as possible if one were to for example download the entire .com zone file.

Let's say not a registry but anyone. Let's say a registrar or registrant would download the .com zone file, which you can get. Then you were to try to generate from that a list of strings that you think would be in the Clearinghouse.

You could fire enough registration requests in a period of time to basically be able to get all of the claims information associated with every name.

Because the system as it works, and this is by the very nature of the trademark claims system and service, by its nature, if someone applies for a string that matches a claim, they're going to get all the claims information.

That's the purpose of it. And the registrant is supposed to say, "Do I want to proceed or not based on that information?" So no matter what level of encryption you have in transmitting that data to the registry, if someone were to apply for every name in the .com zone file, just to see and then not proceed because all it wants is the data.

It will generate most if not all of the files from the Clearinghouse and the information associated with it. So it really doesn't matter what if any encryption you put to it.

CHRIS WRIGHT: Yeah, that's probably one of the fundamental reasons for us proposing the centralized model. If you imagine a decentralized model where that data is distributed amongst all the registries, it will be extremely difficult for those registries to distinguish between legitimate attempts at registering domain names and attempts for people to try to mine data out of the system.

Especially when those requests aren't coming directly to the registry, they're coming proxy by the registrar, potentially even proxy via reseller to a registrar into the system.

In our proposed model where that data comes from the one single source of truth, the Trademark Clearinghouse, the Trademark Clearinghouse has access to all the information every time one of these requests for information is being made.

They're now in the best position to protect that data, to rate limit, to block, to blacklist, or whatever mechanisms we want them to use to be able to protect that data.

That's one of the fundamental reasons why in our proposal we went with a centralized model. I just want to add there are really three things we're trying to look at here.

One is the availability of the service. The availability of the service is very important. As Geoff mentioned before, when the service is down you can't take registrations. Depending on the mitigations you have in place, you may be able to take certain registrations, not other ones.

The second one is the integrity of the system. We need to make sure that the data that we have is up to date and so forth. It goes to the

point that was raised at the microphone before about the data getting out of synch and so forth.

The third one is the protection of the data, the security of the data. That comes down to us being able to detect the data mining and the other bits and pieces that go on.

Unfortunately distributing the database gives us great availability but threatens the integrity and the security of the data. Centralizing the model helps us protect the security of the data and gives us great integrity but it has a threat to availability.

As a community we need to decide what's more important. That's what this really fundamentally is all about, what's more important? And as Geoff said before, counter intuitively, we the registries are saying, "Screw availability. Let's go for the other two."

You kind of would have thought that we would have been the other way but we want to know what you guys think.

KAREN LENTZ: Thanks Chris. So we'll go to the speakers here and then I want to make sure we have time to talk about next steps. Wendy.

WENDY SELTZER: Thanks, Wendy Seltzer and I can keep it short because you, Geoff and Chris, said most of what I was going to say. In asking what is the threat model against which encryption is useful?

And if that can't be described crisply, encryption can do lots of things but it can't prevent people from asking an Oracle that is designed to respond, "Yes I have this name." whether the name is in there.

Thereby getting back an answer to the question, "Is the name in there?" all the layers of encryption won't simultaneously let you get answers when you're the right kind of person and no answer when you're the wrong kind of person. The only way to do that is through contracts or non-technical means.

KAREN LENTZ: Yep, thanks. Thanks, Wendy. Steve?

STEVE DELBIANCO: Steve DelBianco with Business Constituency whose position has been that the Trademark Clearinghouse and other mechanisms, including Sunrise Registration claims notices and URS be as centralized to the greatest extent possible.

Not just for security and integrity, but also for the adaptability and evolution of those services to quickly adapt to new threats. Without having to distribute and coax hundreds of parties to change their code processes and rules.

Now having said all that, I heard Chris talk about a trade-off between availably and security and integrity in a centralized model. We're ICANN and Fati told all of us on that stage this morning that the number two priority frame was operational excellence with respect to supporting its

contract parties, the registrars and registries, in the fulfillment of ICANN's mission and policies.

And all ICANN ever does is registrations and resolutions. And about them all ICANN ever cares about is availability and integrity. So this is time for us to step up, be an adult, and maintain a system that has the kind of availability that we expect out of registry operators.

Look at what Geoff, ask VeriSign, ask Chris about the availability of the zone files today when ISPs are looking for resolutions. That's the kind of availability ICANN delivers.

You have the people in this room to show you how to deliver it. ICANN needs to sign up for operational excellence and deliver five knives of availability with a centralized model for both.

KAREN LENTZ: Thanks Steve. Next, over here.

JAY WESTERDAL: Yes, I'm kind of confused why we're pretending that this information is private. It's public trademarks and you can query for them. Why are we pretending?

GEOFF BICKERS: So Jay why don't you introduce who you are first.

JAY WESTERDAL: I'm Jay Westerdal.

GEOFF BICKERS: Thanks. This is something; it's actually a good question. The question was addressed months ago. We were told from the intellectual product community that although each piece of data, with each single trademark is in itself truly, you could get that from a public source.

It is the collection or the aggregation of all the marks or of maybe one party's marks, like if you have all of, I'm looking at J. Scott right here. If you have all of Yahoo!'s marks, that information in aggregate may provide some competitive or some sort of information that you would not necessarily.

That a. is not easily accessible through any other source in the world and b. may give some information to its competitors or anyone else that it never intended to release.

That's what we were told and I don't believe anyone has backed off of that position. But if in fact that is not true, then we want to hear that as well.

JAY WESTERDAL: It's impossible to protect this data. It's going to get out there. It's query-able. Day one someone is going to have all the information and they're going to make it available. I don't know why we're protecting it and pretending that it's private. It's as public as Lumfile.

KAREN LENTZ: Thanks, Jay. And over here, if you could please state your name for the transcript.

MIKE O'CONNELL: Mike O'Connell from (Inaudible).

KAREN LENTZ: Can you move up?

MIKE O'CONNELL: My apologies if I missed your earlier presentation. I wasn't around. Much of what's been done in on the SSL side on HTTPS and certificate authorities could be used here in the trademark where the Trademark Clearinghouse could sign a request.

I'm not sure if that's what the SMD does but then the trademark holder could then hold then and present that when you distribute your certificate authority.

GEOFF BICKERS: Yep, you've described exactly what our model is.

MIKE O'CONNELL: So the discussion about the alternative claims, is that in conjunction with?

GEOFF BICKERS: Yeah, no, so that's Sunrise and what you're talking about works perfectly for Sunrise. Claims is completely different and any way we can apply the SSL certs and principles and stuff to the claims model. The claims model is very different.

MIKE O'CONNELL: I see. Thanks.

KAREN LENTZ: Thank you. Thomas?

THOMAS NARTEN: Yeah, so Thomas Narten. I'm the liaison to the board from the ITF side. It's interesting sitting through here because it seems like there's a technical discussion going here to a large extent which is not usually what I'm used to hearing.

But a lot of the points that I was going to make actually have been said. Wendy did actually a pretty good job of covering the bases and what I just want to reemphasize is that I agree with the previous speaker here.

I think it's Geoff that this data is not going to remain private. I would encourage people to talk to security experts, people that understand this kind of stuff because if you can query the database.

And Geoff, you said how you'll do it. Get the second level domains out, dot com, dot net, dot org, feed them into the system and you've already got a nice list of all the marks that are interesting. This is not hard to do.

You can say well, if we centralize it we might be able to analyze queries and rate limit. No, this is an offense/defense. If there's demand for this information, the offense is going to win and they only have to win once. You're presumably worried about the bad actors here. And most registries are good actors.

Technical means aren't necessarily the best way to deal with this. Contracts and compliance and things are better tools because they get at the heart of the problem as opposed to relying on technology to do it.

KAREN LENTZ: Thank you.

GEOFF BICKERS: Yeah, I hear what you're saying. I think contracts are not going to help. I mean you're basically putting all the burden, if you were to tell us and we were to all concede that point. Then what I would say is to the registries you don't need to worry about it.

You will have the decentralized model. You don't even worry about protecting the data. You don't even put any mitigation, anything in there. Because right now under the decentralized model is we're going to give all the data to the registries.

We're going to contractually try to bind them to make sure that they put everything in place as possible to try to prevent that data getting out there.

I'm worried as a registry that I have to put in a lot of money to put systems in place to protect that data, to secure it, to do as best as I can to keep that data from getting out.

Because if I'm the one source where that data gets out, this community is going to come after me, maybe by contract, maybe by whatever other means.

I don't want to be responsible for that. Nor do any of the registries. I don't want to sign a contract that's going to basically try to bind me to things that you've just admitted is something I'm never going to be able to protect against.

THOMAS NARTEN: So maybe my push back is not really to the registry. It's to the requirement if it comes to the intellectual property community that what they're asking for might be something they might want to have.

But It's not likely to play out that way anyway. And they should just be cognizant of that. And the other, just a sort of a side note, I find it surprising that you don't have important, critical, private information that you don't protect and that you aren't comfortable that you're protecting adequately. That's a general statement to all registries.

GEOFF BICKERS: We absolutely have that for the data that we've designed our system for. Now you're telling us to each build other separate systems that are one time systems that are, for lack of a better term, throw away systems.

THOMAS NARTEN: Well, but these are solvable problems in general. Businesses do this all the time is my impression. Let's not have this conversation.

GEOFF BICKERS: You know that's absolutely right and we could and we would do it but then again we're also trying to meet the goal of not charging intellectual property owners and registrants additional money for building the system that we have to build for this process.

Right now under the model the Clearinghouse is collecting, the Trademark Clearinghouse is collecting money from trademark owners for filing and registries. But now we're talking about if we need to build what you want, and we can do that.

We've done that with many different things. The intellectual property community and others shouldn't come back to us and say, "Hey, you're now adding charges on to it." we're trying to solve a bunch of different goals here.

THOMAS NARTEN: Right but you're also then shifting costs back to the centralized model and they will have to have sufficient resources and redundancy and reliability. That has cost too. Maybe it'll still be cheaper overall that way, but you're pushing cost somewhere else and there aren't any.

GEOFF BICKERS: Yeah you're pushing it to one as opposed to pushing it out to another. Absolutely right and there are trade-offs. That's the kind of things we have to look into.

THOMAS NARTEN: But again, the big point I think is to really step back and understand whether the data can fundamentally be protected given that it's public.

And given that there will be ways to query it and get yes/no answers and get the information for where there's a yes.

And for actors that really want this information, it doesn't seem like it'll be very difficult to pull it out. Once one of them has it, presumably they all have it.

KAREN LENTZ: Thank you Thomas. We're getting a long line here. So can we have the next speaker please?

SCOTT AUSTIN: Hi I'm Scott Austin. I'm a trademark lawyer with Gordon and Rees. My affiliation is with IPC. My question has to do with the database itself.

Is it dynamic and if so, by that mean if a mark is submitted one day or registered one day and the next day it's abandoned, cancelled, becomes subject to some kind of a proceeding, how soon will that show up in the database?

Or is that for the registrant to challenge? And how does that dynamic aspect in terms of a lead lag affect the centralized versus decentralized?

KAREN LENTZ: Yeah, so I can answer one part of that which is as part of entering data into the Clearinghouse you're agreeing that if it does become abandoned or cancelled or some such that you'll notify, provide notice of that. In terms of how quickly that could permeate, be distributed to

the system, I think it's fairly quickly. If you set a refresh interval and it would always be within that time period.

SCOTT AUSTIN: Does the Clearinghouse check up on that claim? I know they agreed to it but the individual registrant, the person who registers with the Clearinghouse, that's their duty.

KAREN LENTZ: Correct and there's also a dispute like process where if you're aware that somebody is registering names or doing things based on a Clearinghouse record for a mark it's no longer valid. That could be brought to the attention of the Clearinghouse.

SCOTT AUSTIN: Okay.

KAREN LENTZ: Thank you. Cherine?

CHERINE CHALABY: Cherine Chalaby, member of the board. I just want to pick up on a point that Steve DelBianco made about operational excellence. Yes, definitely Fadi said this morning and we are committed to operational excellence.

But we have to be practical and realize that this takes some time. It's not going to be tomorrow. It will take time to achieve that. Now the decision we need to make is a more urgent decision.

It has to happen very soon. Therefore the question I put to the community while you're debating and thinking is it in the global public interest to put ICANN on the critical path of domain name registration?

That's the only question. I don't have an answer to this but I'd like to put it in people's minds. Thank you.

KAREN LENTZ: Thank you. Christina?

CHRISTINA ROSETTE: Yeah, just following up on Thomas' point. Yeah, absolutely, very early on at the last session of the implementation advisory group ICANN's own consultant said, "You're not going to deal with the security issue through encryption. You'll deal with it through contract."

The problem with that is that ICANN took the position at the Prague meeting that it wouldn't amend the registry agreement to add in the necessary language.

And we all know what's going on with the registrar accreditation agreement to add amendments to that. So we're stuck in kind of a vicious cycle.

I mean I think we would certainly be interested in approaching it from a contract perspective but it's been made very clear to us very early on that that was a non-starter.

GEOFF BICKERS: So if I could just, could I just respond quickly to that? But Christina, what we want to know is what would that contract say? You registry will not do what we know if going to happen anyway.

You're going to give liability to registries for something that they can't control. I cannot control whether the database gets mined from different sources because again, as Wendy said and as Tom said, you're getting a legitimate response to a legitimate question.

In essence, you're revealing data that you're supposed to reveal. The very nature of the service is you're going to the registry and you're asking the registry for information that you're supposed to get.

So there's nothing. We've looked at this and I don't mean to sound argumentative. But there's really nothing you could put in a contract that would in any way address this, other than you registry shall not intentionally yourself do it.

But I'm not necessarily worried about registries themselves trying to mine the data. I'm worried about the community and everyone that you can't bind by contract doing that.

CHRISTINA ROSETTE: I am confident that if the IPC was asked to come up with contractual language, we would. We wouldn't like it, but we could come up with it. But my point is that registrars and resellers already use rate limiting for whois.

You could use rate limiting as well. I guess I'm just getting troubled by the fact that we're all sitting in this room saying, "We know this can be

circumvented and bad actors can abuse it. In fact, we're going to tell you how to do it. But now we're not going to do anything to fix it."

For me, that's kind of a much bigger issue. Yeah, if the ask from the IPC is to come up with some suggestions, absolutely. We're happy to try and do that.

GEOFF BICKERS: Right. I guess from a contractual standpoint what we've talked about is it's like, we're entering into a contract with you. I'm requiring you to agree Christina that tomorrow the sun will not come up. Would you sign a contract with me that says that?

CHRISTINA ROSETTE: No because I can't.

GEOFF BICKERS: Because you know it's going to happen.

CHRISTINA ROSETTE: But it's a different, what are the commercially reasonable measures that can be taken, etc. You can do it; it's just whether you're willing to. And the answer we were given in Prague was, "That's not a road we're going to go down."

If that's a road we're willing to consider going down, then that's a way to approach it. But the fact of the matter is having a Clearinghouse up and running with data in it is a critical factor that has to happen before any registry can launch.

So if we're going to go back to a direction that we were all told we were going to discard that needs to happen now, not two weeks from now, and not a month from now.

KAREN LENTZ: Yeah, so Christina I actually don't recall saying in Prague anything about a contractual agreement that wasn't on sort of but we can…

CHRISTINA ROSETTE: There was a rather lengthy conversation. I don't know if he's still in the room, between Marc Trachtenberg and other ICANN representatives at the IPC meeting on this whole issue.

GEOFF BICKERS: I think what Christina is referring to is that there was going to be some Terms of Use that were signed between the registry and the Clearinghouse about, there was a notion of that but I don't necessarily know of a separate terms and conditions or contract between the registry and ICANN. I do recall discussions of sort of a Terms of Use.

CHRISTINA ROSETTE: There was that and there was also a much lengthier, much more detailed discussion in the IPC meeting in Prague. So that's where that came from.

GEOFF BICKERS: Okay, I wouldn't know, could be.

KAREN LENTZ: Okay, thank you. John.

JOHN BERRYHILL: John Berryhill I'm also a trademark attorney and like Jay I have been confused actually for several months about this notion of secret trademark data. This idea that there are these secret trademarks running around and if you infringe them I guess they jump out of the shadows and grab you.

The closest thing to a coherent explanation of this is if somebody got ahold of the Trademark Clearinghouse data then they could determine what your global trademark strategy is, where you're going for trademarks and where you're not. Now that assumes two things.

Number one, it assumes you're throwing your global trademarks in there when you don't need to. If I'm Coca-Cola and I have five hundred trademark registrations around the world, in order to get Coca-Cola into the Trademark Clearinghouse I need to put in one.

So there's no assumption on the part of the data miner that they have a complete data set. So a strategy that's premised on mining the Trademark Clearinghouse data, of which there is no expectation or even sane assumption that anyone would put everything in there, is just cracked.

We've never really had a community discussion around that basic point of this horrible thing that happens if people know you have trademarks.

I will say there are commercial services that you can sign up with that give you access to all kinds of global commercial data.

And they do charge a pretty penny for accessing that data. It wouldn't surprise me that among those trademark representatives we have here, we may have the representatives of those commercial information providers as well. Thank you.

GEOFF BICKERS: So now you see our dilemma as registries trying to implement this. We have some people saying data shouldn't be protected at all. And we have others that say it everything should be protected. We're trying to design a system that satisfies everyone and we're here to tell you we can't do that. But we're willing to try to do the best we can.

THOMAS ROESSLER: Thomas Roessler, tech liaison to the board, about to throw a wet blanket on Geoff. Stop trying because as Wendy and Thomas said earlier, on the technical level what is unclear here is what is the attacker model?

I'm not revealing secrets if I say that Thomas Darden and I spent a little while over lunch at one point looking at this entire thing trying to come up with what is the problem it is trying to solve.

We came up with many of the problems that were mentioned here. For each of those problems we very, very quickly had a very useful attack to get to the data despite the defense that was proposed, and importantly to get to the data cheaply.

In other words, the defenders aren't. Therefore the conversation you need to have in this room is not a conversation about technical measures or the trade-off between technical measures or contracts.

It's in the first place a conversation about what the goals are. We're talking about a service that is supposed to reveal certain information, information is revealed at a rate of so and so many trademark registrations per day.

At what point is the information out there? How expensive is it to gather the information? What is the data worth? From there, I think you can start to have a rational conversation.

That is, first a rational business conversation about what the actual protection goals are. Then a conversation about what the acceptable cost of protection is. Then a conversation about what the technical measures are, that may or may not be a useful part of that protection.

Before you know what that goal is, this is futile. So please don't try addressing something where you don't know the goal.

GEOFF BICKERS: Thanks Thomas. I can't disagree with anything you've said. The only thing I would say to add to it is the protection of the data is only one element of what we're working on with the centralized versus decentralized.

The other elements in a decentralized model are basically relying on a number of different entities that are going to implement the trademark

claims service in potentially different ways. We talked about availability of the Clearinghouse.

One of the things I want to talk about is kind of the reliability of making sure that every registrant is notified in a similar way and to ensure that trademark owners get their information out to the registrant, that the registrar and registry collect the right information to get the acknowledgement.

And we still believe that the centralized model is a much better way to do that. The privacy is one element and we kind of got caught up in it. But the other elements, what's that?

[background conversation]

GEOFF BICKERS: Yes, privacy and security are not the only issues and we got…

[background conversation]

GEOFF BICKERS: Right, there are other issues. There's whether uses can be expanded.

J. SCOTT EVANS: We've gotten off on the security issue because for some people's Red Herring but there are other issues and attributes that you all said. And

Steve talked about those, the ability to adapt the system and only have to adapt in one place so that it's easier to do.

The update issue, so that everyone has the same information available at the same time so you don't have a register in one area that has some sort of natural disaster that can't get the information downloaded to them because the files are too big.

There are all kinds of things. Security is just one issue. Let's not forget that there are other things that you all talked about. Maybe that's the argument to put it out for public comment so that everybody can see all of those things in one place, rather than trying to discuss them linearly. Then we get off track and we don't remember there are other things.

KAREN LENTZ: Thanks J. Scott. Rubens.

RUBENSS KUHL: Just a quick comment that you have 20 new TLDs per week. The claim goes for eight weeks. So an attacker could get 160 attack vectors. If you establish the decentralized model, one has a rate of like one trademark per second. It can amplify that by 160 going through all the registries that are going through claim service at a given point in time, just a quick comment.

CHRIS WRIGHT: So I just wanted to be clear that we didn't come up with this model because of the security perspective. In fact, we started these

conversations telling ICANN and the other people in the room exactly what you guys are telling us, this is futile.

But we were told unequivocally this is what you wanted. So we went and designed a model that gave you what you wanted. But I'm extremely pleased to hear that this is not what you want.

It doesn't actually change our proposal anyway because as has just been discussed, we believe there are other benefits that are worthwhile anyway. But I just wanted to be clear that we're being told that there's this need to protect this data and so forth.

So we've been trying to give you what you want and show you that we've done our best to protect it as best we can. But you know what; it's kind of useless anyway. At the end of the day you all have choice. At the end of the day you're going to choose to put your marks in the Clearinghouse or not.

Just like any other business will do, you're going to make a decision about that. You're going to say by putting my mark in the Clearinghouse I know now that anybody around the world that tries to register a domain name will be displayed my mark information. If you don't want that to happen, then don't put your mark in the Clearinghouse.

That's the whole point of the program. That's what it does. It makes it so that your mark is displayed to someone who tries to register a name before they register it.

If you don't like it, don't put your mark in. that's really simply, that's what the program does. Otherwise if we don't like it, we need to go

back to the start and figure out another way to do this. I don't think anybody wants to do that.

KAREN LENTZ: All right. Thank you, so thanks to everybody who sat in here so late and participated in the discussions. A couple of things on next steps, I have two things in my mind here.

One is request to have written documents and things to review for public comment. The other is I actually think having the live discussion has been really useful. We can do sort of webinar, conference call, and discussion type of things.

It doesn't' have to be either or, but quick show of hands, who would be interested in having written documents for public comment and review. Okay, live webinar discussion type things, a few people.

Okay, that's useful feedback. So we'll take that into account. I want to thank Chris and Geoff for spending so much time up here and answering questions. We will close for the evening. Have a great night.

[End of Transcript]