DUBLIN – How It Works: Domain Name Registry Protocols
Sunday, October 18, 2015 – 14:00 to 15:30 IST
ICANN54 | Dublin, Ireland

DAVID CONRAD: Hi, I'm David Conrad, the ICANN CTO, and this is the third session in the How It Works tutorials. This one is on… What's your title? Your talk, not your title, specifically.

ED LEWIS: This is Registry Protocols.

DAVID CONRAD: Okay. So the session is on registry protocols. These How It Works tutorials originated at the Buenos Aires meeting, the last meeting that ICANN held. The intent here is to provide the ICANN community with more knowledge about the protocols and technologies that we use here. This session will be given by Ed Lewis who is one of the members of the office of the CTO. He actually reports to me, so if he messes up, I'll be able to beat him over the head immediately. If anyone has any questions…

ED LEWIS: It wouldn't be the first time.

DAVID CONRAD: Wouldn't be the first time, yeah. Please direct them at Ed. And with that, I will hand it over to Ed for this session. Thank you.

ED LEWIS: This "How It Works" is about the protocols and the registry [speaks]. But to give a good introduction of that, I'll talk about [inaudible] in general. Then we'll get into what the protocols are going to be. Can we go to the next slide? First [inaudible] ourselves. Next slide.

I like to start with going to the dictionary. We start out with [inaudible]. What's a registry? A registry is just a place that lets you register [inaudible] list of whatever you want to do. Merriam-Webster has official records are kept or it's a [inaudible] official list.

First we're going to narrow it down into something called domain name registries. The thing we're going to keep track of domain names. There are many other kinds of things [inaudible] registry, but domain names are the ones we will look at.

Now, domain name registry is going to be a registry of domain names and how they are put to use and so on, [inaudible]. This tutorial, I'm going to concentrate on just one kind of domain

name registry, which is probably the one that is most visible. That's for a top-level domain name registry.

Now, other kinds of registries to keep in mind. If you start thinking about what a registry does, there are RIRs – the Regional Internet Registries. They keep track of Internet numbers, the addresses, the routing system, the routing parameters. There are the protocol parameter registries that IANA is running. They're just this number this means in this field of this protocol and there's like 2,000 of those plus from the previous tutorial.

There are also other kinds of registries out there that have similar functions – things like land. Who owns land in a country? Motor vehicles, who owns that car? License plates, driver's license and all that is all [inaudible]. Gift registries, who gets to give what to somebody? All of these things are just a list of here's an object and here's what we're doing with that object. That's all this is all about in registries. Next slide.

Now, in the DNS tree, there are different kinds of registries to keep in mind. First, there's the root – the root DNS. That's the registry that's being managed in IANA that has the list of all the TLDs out there that you hear so much about at ICANN meetings.

One level down we have the TLDs themselves. They're broken into different kinds of TLDs. All these TLDs do essentially the

same thing. A domain name belongs to somebody. They're responsible for that domain name and here's [inaudible] information is.

I have gTLD, ccTLD, IDN ccTLD, and other TLD listed up there. G is generic TLDs. Those are .com, .net, .org, and so on. ccTLDs belong to country codes (.ie, for example). IDN ccTLDs are the strings written in non-ASCII writing systems that represent country codes out there. And there's other TLDs. There's just other ones that are harder to classify at this point. They've been around forever or whatever the reasons out there.

All those TLDs below them generally just have registrations to private organizations that get to use the Internet for [inaudible] or not. So let's go to the next slide.

Service level definition of a top-level domain. Next.

The reason why anyone uses a protocol in life is to communicate with somebody else. What I want to do here is take a look at who do TLDs talk to? What are the [inaudible] bodies they talk to, so you can understand why they have a protocol, what the protocol is trying to accomplish. What's the reason for talking in the protocol's definitions?

In this first slide here, I have a picture of the TLD registry – the big can there is a big database of things. In front of it, it has I think – let's start with registrants.

The registrants are names given to those people who own or want or allocate a name. They're companies or individuals that get a name in a TLD. Now, they may go through a reseller, which is someone who just basically retails out those names for the amount of time that they get the name.

Below that, the next step closer is the registrar. The registrar are the companies – you hear a lot about it at ICANN also – whose job is to retail names out from the one wholesale TLD out there.

TLDs are pretty much run as a central area of allocating names [inaudible] referee for who gets a name and puts it in the database. Lots of rules for them. To allow this to be a more competitive environment, we have many registrars that will talk to the TLD to say, "Here's a name I just sold. Please put it in."

One thing I will say that's not on this slide that would be on some people's slides, and probably should be on ours too, is the role of the DNS operator.

The DNS operator is another supplier to the registrant that may have an impact on the system here, but we don't have that represented here, because in ICANN, we have [inaudible] DNS

operator to the TLD registry interface. That's a bit broke down a little bit easier to get between the two, but I want to point out that we don't have that in this slide. Next slide.

Now, another side – it's not just the other side. There's another side to this. These are the organizations that are involved with making sure that TLDs are able to be a reliable, trustworthy place to keep track of things.

If everybody goes to Steve over there – we're going to pick on Steve a lot today – and I tell Steve, "Remember this for me," and Steve goes out and has lunch and sees Dr. Jameson and Dr. Guiness, and doesn't come back, we lost everything. So we want to have someone who's backing Steve up in this game here, and this is what this slide's about.

TLD registry operator, while they're working trying to remember what they've been told by people on the previous slide, here they're trying to store in things like data escrow and also doing some trademark protection of this stuff on the other side. These are support [inaudible] TLDs that are being [inaudible] out there.

This is another important part of TLD registry operations which doesn't get a lot of limelight, a lot of exposure to it, but it's very necessary to making this a reliable system. Next slide.

Of course, the TLD registries also talk to everybody. They touch the general Internet population. When a name is put into a registry, it's registered there. That means that's now I know who it belongs to. But that's now useful to the general Internet, until they're being told "this name, go there." If I want to look up a website, I have to know where to go. So the TLD registry actually has an entire part of – probably the largest [inaudible] TLD operations is running the servers that let people know what it's keeping track of, [inaudible] DNS, the domain name system, and also the WHOIS system. We want to find out who's knocking on my door.

Now, [in a talk], after this, it's going to be kind of a laundry list talk where I go through a bunch of different things, and they're laid out here. These are the six protocols that I'm going to be talking about.

First, just to go around, the EPP (Extensible Provisioning Protocol). There's DNS and DNSSEC are two more. WHOIS protocol, RDAP protocol, the data escrow protocol, and the trademark clearinghouse. Those are going to be coming up not in any order that you see on the slide. It's just drawn out this way.

First, let's talk about the DNS name system. Go to the next slide. Most likely coming to ICANN, you've probably at least heard of

heard of DNS. What it is… It's not a search system. It's a lookup system. I have something I want to know about. I want to know about this name. What do I want to know about it? I want to know maybe the address. What's the address of the website at this name? I look things up [inaudible] information, and the response is simple. It's either here's the address you want, the information you asked for, or no. It doesn't exist. It's not there. Just no. Next slide.

It's one of the earliest protocols on the Internet. It's not the earliest. It's one of the early ones.

It comes from a very long time ago, which is an impact on the way it's put together and the way we try to improve the protocol itself. It's proven to be such an integral part of the system that we just can't uproot it that easily, although many people would like to at times.

The entire realm of domain name registries, the entire industry that we're talking about, just because of the DNS – trying to [manage] what goes in and out of the DNS. What the DNS means to a registry, it's the most visible part. Being up or down, basically if their DNS is running or not.

The DNS is probably not the most central part of a registry's functions, but it is the most visible piece, the highest profile

protocol we have, and it [approaches] being an element in the way things are running out there.

It's by far the most used protocol out of the TLD. It's untold numbers of people that rely on this. We don't really have a roster of all the people who ask questions. We can tell, but even that's filtered by the [protocols put together]. And the [inaudible] queries are anonymous.

In a [inaudible] diagram, the orange bubbles here are things external to a TLD. The blue are inside the TLD. An agent of the registrar will register information with the registration site which filters to the database down to the DNS server. That DNS server is [now] going to speak [to the] DNS protocol.

IANA is up there also because IANA is involved with the setup of this, [in the sense] that when a TLD gets granted and put the root zone, the name server setup, [inaudible].

[inaudible], there's authorative servers, and that's what the TLD is [inaudible] servers. There are also recursive servers which are essentially middle boxes inside the system. Then there are [subs] and clients out there which actually are doing the [inaudible].

And if you stack this up, you see the DNS server from the previous slide speaks to the DNS protocol. It speaks to recursive

servers generally, but [inaudible] speak the same protocol – the same DNS protocol – out to whoever's asking for this information. The DNS protocol is much larger than just the TLD operator.

The next set of slides is going to be DNSSEC. Now, what the previous graphic showed was that there's intermediaries between the blue bubble on top, which TLDs run, and where people.

For this reason, the DNS has to be secure. In a lot of places, when you are on the website, [you can] go right to the website and use a point end-to-end security by encrypting [inaudible]. DNS doesn't have that. We don't have an end-to-end connection between those using the information and those [inaudible].

So DNSSEC was something that had to come up to help protect the data as it flowed through the DNS system to the end point. DNS data is transferred in the open, so we had to attach some way of writing like a digital check or a digital signature so people could verify that they got the right stuff.

End-to-end encryption like HTTPS isn't the solution here because you don't have connections all the way through the system, [inaudible]. The history of DNSSEC, it was developed in the 90s. [inaudible] workshop with operators through about

2004 and that's when the first documents came out of the IETF [inaudible] the protocol.

It hung around for a couple years as [inaudible]. Didn't get a lot of attention, until a researcher named Dan Kaminsky [inaudible] called the end of the cache as we know it. It basically said that there are so many ways to abuse the fact that DNS is not end-to-end, [inaudible] pieces. We need to do something about this. That's when DNSSEC kind of [appeared].

In 2009, it appeared in the first TLD out there. Sweden was the first to sign their zone. And by 2010, the root zone was also using DNSSEC. Next slide.

[inaudible] approach is the data that is being transferred is a company by a digital signature. To validate that, you need a public key. The public key is going to be available through the DNS and it must also be trusted. This is all based on public key cryptography, and the cryptography boils down to I have a pair of keys that work in conjunction with each other. If I encrypt something with one key, the [inaudible] back and forth. So these two are tied together mathematically. [inaudible] one of those keys and I gave it to everybody, that's the public key. I take the other key and I make it a private key. The idea is that if I sign it with a private key and you can validate it with a public key, you

know I had to had signed it with the key that's private. [I'm the only one] who has that private key.

That, plus chaining these keys together. We've made a scalable, trustable network. A framework for this protocol. The hierarchy that's used is matching the DNS tree that we're protecting.

What does a registry do? The registry has to manage its own keys. It's in charge of keys that it uses. It has to be able to accept registrants, something called a DS record. [inaudible] DNSSEC record that helps chain these keys back and forth. The TLD signs these records. It [inaudible] as if they're their own.

The other big part of this is signing the negative answers. It sounds kind of weird to mention that as a primary piece of work, but signing negative answers is a lot of work. A lot of time has gone into saying no, because basically in DNS, I would say no to you by saying, "Here's an empty answer. There's nothing to see." And you can't sign nothing. You have to sign something. You have to say something to sign it.

Finally, the last part is interacting with IANA at all times to make sure that their DNSSEC data – in fact, their DS record for TLD – [is correct] in the root zone. So they have to now interact more with IANA than they used to.

And the [block] diagrams here, they're not exactly the same order, but again the registrar through the agent of the registrant registers names – registers actually DNSSEC data through the registry database to the DNS server. Now it becomes DNSSEC.

DNSSEC isn't a completely different protocol in DNS. It's an addition of some records, same server set and the same network.

This week there are two sessions that are geared to DNSSEC. One is a beginner's guide for everybody. That's actually going to be in about a couple hours from now. I'm sorry, it's tomorrow, Monday, from 5:30 to 7:00. On Wednesday, there's an almost day-long workshop on DNS with a lot of presentations.

By the way, if there are any questions about this material, just stop me. Go ahead. We got remote people I think.

BRIAN AITCHISON: Thanks. My name is Brian Aitchison. I'm with ICANN Operations. You talked about DNSSEC and this validation key. Is this something that I as an end user of the Internet would see or that I would know about? How does that relate to my…?

ED LEWIS: Sure. Since I was talking about TLDs, I kind of glossed over that. As someone who's using DNSSEC, someone receiving the data, in order to do the validation of DNSSEC, you would have to get a copy of the public key that ICANN has put out for [inaudible]. There are other options, but basically you'd want to have the key that ICANN is producing for this whole process, which is actually not something that would be related to TLD operations.

The next protocol I'm going to cover is WHOIS. WHOIS is ancient also. It predates DNS. It was around before that. The job of WHOIS is basically who is the person who has this IP address? Who is the person that has the domain name?

It's a way of identifying the operator at the other side of the Internet. That's the purpose of WHOIS. It's very simple. Simplistic question and answer protocol. It came up at a time when this was all research oriented. Everyone knew about what was going on out there, and basically it was essential to keep the network together to know everything [inaudible].

WHOIS protocol [definition] is pretty much this entire slide. It's just on this slide. DNS is a very complex protocol. This is a very simple one. Open a connection to a certain place. You send a question. You receive an answer and you close the connection.

What hat question is and what the answer is never really defined. It's just whatever it is is defined by whoever [inaudible]

techs going back and forth. [inaudible] TLD, it's registry database feeds the WHOIS server which speaks the WHOIS protocol out to clients. The orange clients are being run by anybody out there. The blue [inaudible].

Now, where do we start with the problems of WHOIS? [inaudible] WHOIS is kind of a topic of discussion at an ICANN meeting.

The first thing is this freeform-ness. When you have [inaudible] question, just to get an answer back, that's not really helpful to inter-operability. When you're trying to write software that uses this stuff, whatever the worst thing to try to do in [inaudible]. Just whatever. Just write it. It's very hard to do that.

Second thing out there, too, is that the original spec basically assumed ASCII. Everything was he 26 letters upper and lower-case, whatnot. There was no consideration for international names or anything non-US names at the time.

For those who are writing [inaudible], there were no other kinds of answers. No other… We call them meta-answers. There's no way to redirect people to another place. WHOIS assumed everything was in database. That's kind of a bad idea for [inaudible]. We couldn't say, "I don't have answer. Go somewhere else." DNSSEC had that, but not the WHOIS protocol.

Differentiated access was impossible because you didn't know… Everyone was [inaudible]. Whether or not you want to do [differentiated] access is one issue, but [inaudible] WHOIS protocol is today [inaudible] do it.

And finally, there is no way to validate the data in the answers. Because everything was freeform, not even postal addresses had to be in a form that was verifiable. You couldn't tell. Next slide.

There are two sessions here for WHOIS in the week. The WHOIS Review Team International Data Expert Working Group is going to meet. You've got to go fast for that. Second is the Thick WHOIS Policy Implementation will be on Wednesday morning from 8:00 to 9:15. That's another place where they're talking about how to redo WHOIS the way it is today. [inaudible] WHOIS.

EPP is Extensible Provisioning Protocol. I'm sorry? Thick? In the world of the way things have laid out, there are two kinds of models for registering… Actually, could you go back to the third or fourth… The slide that showed registrants way out in the beginning?

When the registrant comes and says, "I want a name," [in order to] resell to a registrar, the registrar will say, "My customer is so-and-so. They want to buy this name." Now, they have to turn around and tell the TLD that this name is now going to be used.

The question here is thin versus thick [where] it comes into. If the registry is thick, it has all the information the registrar has. The registrar says, “Here’s this domain name. Here’s where it’s going to be located [in the] name server. Here is the person. Here’s their admin contact. Here’s all the information that belongs [inaudible].” They have all the information, thick.

If it’s a thin registry out there, the registrar only tells the [inaudible] the domain name and here are the name servers. Then anything else you want to know, come back to me and I’ll fill in the rest. Thick and thin refers to what does the registry itself contain. You can flip back to where we were.

Thick WHOIS policy, that was saying that they want to take all the thin registries out there and make them act like thick. WHOIS [inaudible]. It’s kind of complex. One of the issues in WHOIS is that the word WHOIS has gone from being just a protocol to being a euphemism for the entire registration system.

They’re trying to make the WHOIS protocol act as if everything was [inaudible] WHOIS protocol purposes.

EPP is business-to-business protocol, meaning that it goes between companies that have a pre-arranged agreement. WHOIS and DNS [inaudible] anonymous protocols out there. Anyone could ask the question. The EPP, you have to have an agreement between the two sides. The purpose of EPP is to edit

the registration database. [inaudible] modify domain names, contact information, transfer, and just about anything else that involves the business of registering names.

This was built about 15 years ago. I think I have [inaudible] on the history of this. It was taken from experience with other registration systems and [inaudible] pretty much early on from some ICANN work. In 2000-2003, it was developed in the IETF as a working group. It was based on the earlier protocols that came from the common net experience. There was an RRP at the time that's been modified. It's actually pretty much gone away and EPP has taken over for it.

In the IETF, the next about six years became a full standard which is kind of rare in the [inaudible] actually became something and finished the entire process of reviews.

It has been mandated for all the gTLDs and sTLDs, [the] ones that ICANN oversees. I'm not sure if it's total or most at this point. I'm not sure if everything's come under agreement. It's hard to say because the agreements apply until they wear out. [We're] not sure where some of the other TLDs, where they stand.

As far as ccTLDs, they have not been required… ccTLDs of course are not under the review of [inaudible]. Nevertheless, a lot of the ccTLDs as this went on that this was a pretty good idea.

The protocol actually worked pretty well and they adopted this protocol in their [inaudible].

Currently, in the IETF, there is a new working group which is managing the extension [part of this]. The "E" in EPP is Extensible. Some people have been mad that it's been too extensible. Many people have written [inaudible] to do the same thing many different ways, so we're trying to come back down to let's just have fewer extensions to manage right now.

The EPP does not do an exclusive piece of work. It's one way to cause [edits] to happen to the registration database. It doesn't mean you can't have something else. I've worked at registries that do have other ways of doing this. [inaudible] to other things.

Now, there's a policy out there. Policies are all over the place here in registries. They run by certain rules [inaudible]. That may be [inaudible] whether they use EPP or not. It's not a technical barrier.

Now, the EPP protocol uses TLS or a strongly secured transport layer. We actually spent time looking at what's underneath the protocol. It's going to be a strong [inaudible] area. It's coded in XML because that was the way things were done back in the day it was defined. ML is extensible [markup] language. It's probably not the most [advanced] today, but it is what it is then.

The server sits inside the registry, and again, the diagram down below shows it. The client of EPP belongs to the registrars, and generally the registrars will run their clients and they will send the queries in through EPP to the EPP server [inaudible] ask for a domain name to be registered or to modify a domain name. [inaudible] EPP.

RDAP. The registration data access protocol. This is basically a new version of WHOIS. [inaudible] response protocol to inspect the registration's database. It's basically solved a lot of the problems that WHOIS had. It's just been defined. It was [invented] by the IETF in the last few months. It's build on top of [inaudible] new protocols. [inaudible] a lot of web developed generally. Next slide.

[inaudible] server, which parses the queries. It looks information up in the database and it prepares a response. [inaudible] is a web browser, API. There's a lot of [abilities] built into it and it can perform authentication steps.

The history is that there was dissatisfaction with WHOIS amongst the RIRs, and so a couple of the RIRs went off and decided they were going to look for a better way to report on IP addresses and other routing parameters. They had a very successful experiment with a web-based approach, which is [inaudible].

From that [inaudible], RDAP is very much closely tied to being [inaudible] WHOIS and very closely tied to the way the HTTPS [inaudible]. One other flaw out there was that it really drove the commonality, the common elements of registering names and numbers. [inaudible] not that different between the two sides of the [query], the numbers side and the names side or the policies are quite a bit different [for either] side.

The query looks a lot like a URL. It's a formalized version of a question [inaudible] back over HTTPS. It uses JSON (Java Script Object Notation). It's a formalized saying what the answer is. And it lets people now [inaudible] the answer back in a programmatic way. With WHOIS, you had to read it on the screen or scrape it and decide what it meant. This actually has a formal [inaudible].

There's also a way to redirect someone saying, "Don't ask me. Ask somebody else. I may not have that name in my area, but I know who would have that for you somewhere else."

What's remaining now, this protocol has just been created out there. In a lot of places, an operational profile is being defined for how do we actually take the specification and make [inaudible] thing. Next.

It has a defined data model. It's expansion-friendly with the queries in [inaudible] format, so we know how to add things to

that. It can go beyond ASCII. That's one of the [key] things we wanted to do.

We can distribute the data sources. It's not just one database for all this. It's not [assumed] to be one database. We can have differentiated access, meaning that we can give people passwords to certain parts of it, [inaudible] what that model is going to be. You don't have to have that, but you have the option of having hat now with RDAP. It's a more modern protocol. It's more compatible with the current way software engineering is done.

And here, RDAP client speaks to the RDAP protocol to the message handler, which goes to the database. Actually, before you go, the arrow there is intentionally going from the database to the message handler. There's not a way to… It's not an update protocol. It's a read-only protocol.

This week on Wednesday morning there's a session on the RDAP implementation and they're discussing an operational profile as being… I believe [inaudible] the operational profile for RDAP or for ICANN's purposes.

Data escrow. Data escrow is a way to take what's in the register of domain names and [who is] responsible for them, store them somewhere at a third party's location in a way that's safe, that everyone agrees is safe.

Why would you do this? Operator failure. If a TLD goes out of business suddenly, is shut down, it fails for some other reason, we have another copy of it somewhere else.

We eventually could allow the TLD to be restarted by somebody else with this information [and] assume it all.

The storage [third-party] has strict rules for access. It's not just a [inaudible]. It's actually a regulated way of keeping track of this data. It's put somewhere. There are certain rules under which the data can be retrieved.

Now, the history. There was an IETF session where tried to see if the IETF was [inaudible] talking about this protocol. It was in 2010 or so, give or take a year. The IETF was not interested in this.

This doesn't mean the protocol is unimportant. It just meant that the IETF wasn't interested in looking at this, because frankly, it's kind of a simplistic thing to do. How do I just copy what's in my database to a third party and deal with it? It's very simple, but it's very important and very specific. [inaudible] really have a lot of general ideas that need to be discussed in the IETF.

The result. The data escrow definition exists in two places, or it's in two parts. One is in a framework which is in a document that's

not quite an IETF RFC. I think they're kind of working trying to get that done anyway as an individual [inaudible].

The timing of this, which is entirely different issue, which is when do you do this, is in specification 2 of the registry agreements out there [inaudible] for the new TLDs.

What it is, basically, is a dump of the database XML version in one or more files is most likely there. It's compressed and encrypted to keep the cost for storage down. Every day a deposit is made and every week there's a full dump of what's in there and so on, [inaudible] that the registries are able to make these deposits every day. It's important because if they go down, we don't want to go back days to see what was the last [inaudible].

So this shows a [block] diagram with the registry operator on the left. [inaudible] public keys because everything is encrypted. A full deposit, an incremental deposit go across, and the notification saying we made the deposit. And to ICANN on the bottom, notification and public keys go over there, but not the deposit itself. We just get… ICANN gets notified that it was made to make sure that it's being made consistently. Also, we have a copy of the public keys in case ICANN has a reason to go and ask for that.

I guess actually [inaudible] data escrow. Data escrow is something that I've seen mostly involved with ICANN registries,

but when I presented this last time, some people say other… ccTLDs are also picking up this, too. They realize it's an importance to the data. They want to make sure there's a third party that's also managing copies of these across the board. Not just in the ICANN realm, but just about all realms of the TLDs out there in the Internet.

TMCH, Trademark Clearinghouse. This is another support protocol which has been [inaudible] mostly in the ICANN area, but again, I've also heard that the ideas here have been adopted to some other country codes too. [inaudible] but mostly specific to ICANN mechanism for recognizing when a trademark term is being registered in [inaudible] name.whatever is what they're looking for here.

In this few slides I have here, I guess the registry touching part of the protocol – because this actually involves a trademark clearinghouse. It involves other elements out there. I'm going to stick to what does the registry have to do here to keep it simple.

There are two phases that are involved in [a TLD]. When a TLD is created [from nothing], they go through a couple of things like sunrise and trademark [claims]. Sunrise is what do you do when you begin operations. Then you do [inaudible] time where I'm going to let the trademark holders know that I have a new TLD out there. And there are other phases out there like [land rush]

and special [inaudible]. You may hear [inaudible] maybe specific to TLDs out there. The way they start opening names before, they go to general availability which is when anyone can get whatever they want [inaudible] rules they set out.

The two phases involved here are sunrise and trademark claims. And everything here is a protocol built over the secured web. So sunrise refers to the opening of the TLD to trademark holders first. When a TLD goes out there – and a TLD is going to allow others to register names in them (not all of them do) – they have to have a period of time where they let trademark holders come and [inaudible] names first.

At that point, the registry the supplies the TMCH, the list of domain names that are being registered, so that [inaudible] trademark clearinghouse. And coming back, the registry will get a list of trademarks that are no longer being reserved. In this case, the trademark holder will talk to the TMCH. These are [inaudible] or not protect my name and then they can retract that.

During trademark claims, things look different. Anytime that someone registers a name during the trademark claims period, there's a notification sent saying that your name may be a trademark or looks like a trademark and [there may] be a concern about that.

[inaudible] the trademark clearing house. The list of domain names that are registered that got the warnings and the registry gets back a list of labels corresponding to other trademarks to watch out for. Basically, this is a warning period. There's not necessarily going to be action taken on this, but [inaudible] someone may say, "That's my trademark. You can't have it."

Again, this slide shows the interaction with the registry. There's sunrise [on top]. There's the revocation list. And then during claims, it becomes a do not… A DNL. I forget now what the L stands for. Do Not List? Because I have list there also. I'm sorry, domain name list. I almost got through the entire deck without screwing up an acronym – domain name list. So the [inaudible] get a different list back.

All along the registry is actually feeding back to the trademark clearing house. These are names that have been effectively allocated. So that's what goes on in this protocol.

This protocol actually… If you go through the Internet draft that's written on this, it's an incredibly well done, I have to say, a really well done ASCII [art] diagram of this, which was amazing, that I could never get back. But anyway…

So we've actually walked through all of them already. EPP, Extensible Provisioning Protocol, the way to register names. DNS and DNSSEC, which are the ways that the Internet knows

who has what. WHOIS is the protocol for seeing what's registered out there. RDAP is the next version of that. It's in use in many places. Actually, it's in use in some places, but [inaudible]. Data escrow is where we copy the data off somewhere else. And trademark clearinghouse, which is where we protect trademarks from being [inaudible].

UNIDENTIFIED MALE: [inaudible] for Domain Cocoon Incorporated. [inaudible] TLDs need to offer premium pricing. Is there any other way to integrate this into a standardized EPP?

ED LEWIS: Off the top of my head, I do not know. The way to find out is to see if anyone has applied to document their extension that they're building. [inaudible] see if they have that. There may or not be an extension listed in this. In general, if there's a specific registry that [inaudible] looking at, they would probably have to ask the registry if they have an extension because no one's required to [inaudible] the extensions they have.

UNIDENTIFIED MALE: [inaudible] and there's two competing standards for the premium price. The version produced by [Gavin Brown] from [inaudible] standards. Both very good. Both do very similar

things. But in general, the premium pricing needs a lot more work.

ED LEWIS: Do you know if [they have been] submitted to the IETF?

UNIDENTIFIED MALE: There has been, but again, both versions have been implemented quite heavily.

UNIDENTIFIED MALE: [inaudible]. The notification that's sent to ICANN [inaudible].

ED LEWIS: I believe it's that it's been sent. Because of the TLD [inaudible] been sent. I don't think they [inaudible] notified by the escrow holder that they've received it. At one time, I was more involved with this and I just don't recall which. I know that checking is done.

GEORGE MINARDOS: George Minardos from dot-build registry. Maybe I was too slow, but it seemed like you went over [inaudible] regarding the data escrow and [inaudible] of which key?

ED LEWIS: [You] get a copy of the public key, and the public key is used to basically [inaudible] data that's been escrowed. In a sense, if ICANN meets the [inaudible] the data from the escrow agent based on the rules, we get [inaudible] that way and make it available [inaudible] read it for whatever purpose we had to read it for.

UNIDENTIFIED MALE: [inaudible] typing online, do you have another question coming up? [inaudible], so that ICANN may verify the [inaudible]?

ED LEWIS: Yeah. I guess I was thinking of the registry's version. If the registrars are submitting to escrows, too, then it would basically from the source of the… From the source of the escrow dump. Again, preparing the slides, I was thinking about registry operations, not necessarily registrar operations.

[DANIEL LEE BANKS]: [Daniel Lee Banks], dot-ky, Cayman Islands. WHOIS data, thick records, do they [inaudible] or the [inaudible] being asked brought up by the law enforcement community around the world? Does a thick…

ED LEWIS: Generally, [inaudible]. It's hard for me to say anything about policy around the world. But if you're someone who needs to know the address of the person who's paying for this name, you would need the thick records. [inaudible] that answers your question. Yeah, the thick has more information. A thin response would be here's a domain name and here are the name servers where it belongs. I need to have them for the DNS to work. And then I would say the name was sold through somebody else. That would be a thin [inaudible] there.

If you ask the registrar in that case, they should have a thin registry set up. The thick registry, I would just go to that one place and get all that together.

UNIDENTIFIED MALE: In addition to that, with the thick WHOIS, thin WHOIS discussion, one of the challenges that the Internet community as a whole [inaudible] data privacy and the use of that data across different legal jurisdictions. One of the big challenges right now with the model is the need to extract the registrant data propagated across national boundaries and that is something that [inaudible] more of an issue quite recently with the EU data protection, the Safe Harbor ruling that just came down.

You can probably anticipate that there's going to be continued discussion on the question of thick versus thin WHOIS in the

future, and all of it can be seen as almost moot because of the development of RDAP and the requirements that are placed [inaudible] gTLD registries to implement RDAP [inaudible] or some number of days after it's standardized within the IETF [inaudible] recently has been.

UNIDENTIFIED MALE: Just to clarify that, that mechanism [inaudible].

UNIDENTIFIED MALE: You mean…

UNIDENTIFIED MALE: Various levels of access, so a law enforcement [inaudible].

UNIDENTIFIED MALE: Exactly, yeah.

ED LEWIS: The requirement to adopt RDAP is not after the standards [inaudible] after ICANN asks for it. In other words, the operational profile going on [inaudible]. After that's done, that's when the clock would start ticking. So [inaudible] have to have that model even defined.

UNIDENTIFIED MALE: [inaudible] from National Internet Exchange of India. I started [inaudible] working group that [inaudible].

ED LEWIS: The RDAP protocol is to find a bunch of documents out there, and they have different queries out there. They have different [things]. But it's also a [inaudible] piece of work, too. Not everything in the [inaudible] will apply to all [deployments], and in many cases, deployments will have more information. [inaudible]. I don't think it goes into the specifics of [inaudible] contact information is in the response [inaudible].

For example, I don't believe it specifies there's a tech contact for a name. You want to have an operator. You want to have a tech contact [for a name]. The protocol doesn't say a tech, but [inaudible] profile. A tech contact is going to have to have. Now, where it says that [inaudible] operational profile. That's basically someone says [inaudible] specification for the protocol, but here is how I'm going to run it.

I know that in, I think, Belgium [inaudible] operational profile for their implementation of RDAP. ICANN has one that's actually in play. I think it's [inaudible] profile has been put out there for review and that's what will be discussed [inaudible] this week.

To give you a reason for this to really be there, the RIRs have [inaudible] data. They manage data differently than the names do [inaudible] same protocol. So that gives you [inaudible].

If you imagine the difference between the names [inaudible], it gives you an idea of the operational [inaudible].

UNIDENTIFIED MALE: So protocol RDAP is not operational [inaudible].

ED LEWIS: I don't know if any domain name TLD that has a server [inaudible]. The reason why I'm looking around is that ICANN is not the only place [inaudible] TLDs. There's ccTLDs out there. RIRs are operating RDAP. [inaudible] the numbers space, rather.

In domain names, the place I know that's most [inaudible] is Belgium, but I may not know of someone ahead of them. But I know Belgium has done a lot of work.

UNIDENTIFIED MALE: I have a clarifying question from online. RDAP uses [440S], or does it use a different port?

ED LEWIS: RDAP uses HTTPS protocol. WHOIS uses port 43, but that's WHOIS. RDAP doesn't have its own port number. [inaudible] HTTPS protocol through its port. Yeah, but I think HTTPS actually has a protocol.

UNIDENTIFIED MALE: Port 43.

ED LEWIS: Yeah, sorry. Yeah.

UNIDENTIFIED MALE: Last question.

UNIDENTIFIED MALE: [inaudible] slide number 9. My question is where does the privacy proxy [inaudible]?

ED LEWIS: Generally, those three functions, which [would] probably be associated with the registrar, I'm guessing, [inaudible] the registrar may say if you register a name with me, I won't let anyone know it's you [inaudible] some other place. [inaudible].

UNIDENTIFIED MALE: [inaudible] privacy and proxy, so this is [inaudible].

ED LEWIS: [inaudible] you or…

UNIDENTIFIED MALE: [inaudible] equivalent.

ED LEWIS: I've actually never had experience with the two of them, so I wouldn't know. I used to run into places where things were hidden, but I never bothered to – I never looked at it from their point of view. I didn't look at how they do the work.

UNIDENTIFIED MALE: [inaudible] proxy is like a legal entity that actually takes your place as the registrar, so if anyone wants to find out about [inaudible] a registrar [inaudible].

UNIDENTIFIED MALE: I believe that we're out of questions, out of slides.

UNIDENTIFIED MALE: I want to thank Ed for the. As some of you have been [inaudible], we're actively seeking feedback for [inaudible], so please, before

you close up your laptop, [inaudible] how this presentation went, how we could do better, if there's other topics you'd like to see, [inaudible] hear about it.

DAVID CONRAD: I just want to thank everyone for attending. We have one more session. This is the session that follows this one, I guess in about half-an-hour.

UNIDENTIFIED MALE: It's 45 minutes, actually.

DAVID CONRAD: 45 minutes until the next session.

UNIDENTIFIED MALE: We've got some time here. An extended tea, maybe?

DAVID CONRAD: Right. The next session is actually a new one that's being done the first time here in Dublin on the root server system. It's being presented by, I guess, Duane Wessels of VeriSign who is a member of the Root Sever System Advisory Committee. Also one of the root server operators. So if that interests you, please come back at 3:30.

UNIDENTIFIED MALE: 3.45.

DAVID CONRAD: 3:45. Thank you very much.

UNIDENTIFIED MALE: And if also you'd like to stay here, we're [inaudible].

DAVID CONRAD: You don't have to leave. Thank you.

UNIDENTIFIED MALE: Thanks.

[END OF TRANSCRIPTION]