LOS ANGELES – IANA Department - Who, What, Why?
Monday, October 13, 2014 – 10:30 to 12:00
ICANN – Los Angeles, USA

ELISE GERICH: Can you all hear me? I can't hear myself. Good. My name is Elise Gerich, and I'm the vice president of the IANA Department at ICANN, and we're glad you've joined us.

The room's kind of set up in an interesting way, so instead of us sitting up at the beginning of the table, we decided we'd sit back here and look at you instead of have you look at our backs.

So hopefully this will work out well.

The reason we're having this session, which we've actually never done before at an ICANN meeting, is to talk more about what does the IANA functions department actually do.

We have a little subtitle up there that says, "Why are we less interesting than you think?"

Every now and then people come up to us and sometimes they really don't know what the IANA Department does, and who we are, and why we do it, and what we do on a day-to-day basis.

So another subtitle could be, "An hour in the life of the IANA Department."

We have a panel of four of us today. I'll introduce them right now. Kim Davies is right here. Raise your hand, Kim. And he'll be talking primarily about the names area.

Naela Sarras, she'll talk primarily about the numbers space. And Michelle Cotton, and she'll talk primarily about the protocol parameters.

I'm going to try to move this. Sorry. I went a little too far too fast.

The laptop? I have to point at the laptop? Good.

So one of the things I wanted to make sure is that you knew who we are. So the IANA Department has 13 people in it right now, and these are photos of everybody. Many of us are in the room today. If you're an IANA Department person, please raise your hand. I won't make you stand up.

Leo, your hand has to go up higher. They can't see you.

So if you have any questions after we've finished our presentation, please feel free to stop and ask any one of us.

When I started at ICANN four and a half years ago, there were nine of us in the department. And if you count quickly, you can see we're now up to 13.

What we've done is we've added two cryptographic key managers, that's for helping with the DNSSEC and all the key security that we have for DNS. We've added a software developer to help Kim do a lot of the tools and the reports that we generate. And if you're a Python developer, talk to Kim. We're looking for somebody right now.

And we've also added an audit specialist, someone who helps run our third-party audits with Price Waterhouse Cooper. So these are the folks that are in the department and we're the ones that do the day-to-day work for the IANA functions.

So that was the who. Now we're getting to the what.

So what does actually the IANA functions -- what are we? What are they?

So we were established in 1998, and be -- under ICANN. Obviously the IANA function started a lot earlier with the days of the ARPANet and things of that nature, but the IANA functions that are operated by ICANN began back in 1998.

And we manage the registries for the unique identifiers. And you'll be learning more about what that actually means. What does it mean for us to manage the registries?

And those primary registries, unique identifier registries, include the numbers, the Internet numbers, the domain names, and the protocol parameters.

And finally, we manage and we administer these registries based on the policies that are done by the communities themselves. So we don't make up what the registry should be. They're actually made up by the various constituencies that were formulated well before the IANA functions got started. So the protocol parameters, the policies that we follow come from the IETF. The numbers, the protocols are global policies that come from the Regional Internet Registries. And then the

names originally were just ccTLD policies, which are ccNSO, and now there's the GNSO and ccNSO that make policies for names.

So these are our primary functions. There are some other things we do as a department, but these are the primary IANA functions.

And what I'd like to do now -- whoops. I keep pointing at this and it's not turning. There.

I'd like to turn the microphone over to Michelle Cotton and she's going to walk you through what the IANA functions department does with regard to protocol parameters. Michelle.

MICHELLE COTTON: Thank you.

ELISE GERICH: And there's the clicker, and it's that one.

MICHELLE COTTON: Hello, everybody. My name is Michelle Cotton and I am the protocol parameters engagement manager, which mostly means working a lot with the IETF community, the IAB, and other standards organizations.

[No Audio]

If online, you can actually go to www.iana.org/protocols. Or if you're on the home page of www.iana.org, you'll see the protocol registries link. If you click there, you'll see the actual listing of all the registries.

Things like media types, port numbers, everything you could possibly think of, we use.

Are we good?

>> No, the Wi-Fi is off-line.

MICHELLE COTTON: I have it in front of me. We're having some Wi-Fi issues it looks like, so I think we're just going to continue on.

Where to protocol parameter registries come from? We've got this whole listing of registries that we maintain.

They mainly come from the IETF community. The IETF community members, which is volunteer members, get together and write Internet drafts, and these documents, they describe various different protocols, applications, all different types of things. And when these documents get approved, they go through a process, an IETF process. They get approved by the leadership of the IETF.

There we go. Thank you.

Okay. Kim, you're controlling now? Okay. (Laughing)

Kim's in control.

When these documents get approved by the leadership of the IETF, they become official, what we call, RFCs, Request for Comments. Request for Comment series was originally developed by Steve Crocker, and it's

been around for quite a long time, and I think we're now in the 7,000s of RFCs. And if you want to know more information about that, you want to go to www.rfc editor.org, and you can learn all about them.

So these RFCs contain instructions, guidance to the IANA Department on what we need to do. So instructions on creating new types of protocol parameter registries. How does a person or a company get a new code in that registry? How can somebody update that registry? What is the registration policy? Can anybody come to the IANA Department and say, "I would like a new number in there"? Do you have to have a technical expert review your request? Do you need to actually write a document, another RFC to get a parameter in that registry?

All those instructions are in these documents for the IANA Department to follow.

So next slide.

So what's our role with these protocol parameter registries and these RFCs? So before the RFC gets approved for publication, we actually do quite a few reviews of these documents to make sure that the language in the RFC is absolutely crystal clear on what we're being asked to do.

Create registries, update existing ones. We're just making sure everything is stated clearly. And sometimes we'll ask the folks writing the documents. "This isn't quite clear. Did you think of this?" Or "You have a small registry. Can you please state the registration procedure?" Maybe they haven't thought of how they want people to get those

parameters and that registry. So we work quite a bit with the authors. We do lots of reviews before the document gets approved.

After the document gets approved, it's time for the IANA department to implement those actions. So what we do is we create the registries. We put in the initial assignments. We put in the reserved values that have been requested. And we add them to our registry with all the information that's helpful to the Internet community on where the registry came from and how you can get values in parameter registrations in those registries.

We also maintain them. We've got quite a few registries that are very, very active. And then there's a lot of registries, many, many, many registries that don't have a lot of activity yet. We still maintain them. And our goal is to keep them public at all times for the community's use.

Next slide. So this is an example of what an IANA consideration section in one of those RFCs looks like. And you can see the three highlighted areas are the main instructions here to the IANA department.

Creating a registry, maintain it, how future registrations are received.

So, in this case, there's a policy called first come, first served, which means you submit a request. And we, basically, assign it or put your name registration in these registries without any technical review or documentation required.

And all the different types of policies that the IETF has for registration procedures are defined in RFC 5226. So, if you're interested in learning about all the different types of policies that they set, go take a look at

that document. It also has a lot of other interesting information on how we review documents and work with the IETF community.

The IANA section also tells us that there's no initial assignments for this particular registry. Often, it will list all the initial assignments that this particular registry is going to put in there just with the initial approval of the RFC.

So, again, we work very closely with the IETF community. We're just making sure that these instructions are clear. And, if you look at any documents at the rfceditor.org Web site you can look up IANA considerations and see lots of examples of the instructions that these documents give us.

Next slide? Okay. So, after RFC approval, we create these registries and update them. We maintain them, follow registration procedures. And often we also update references. As the IETF produces more and more documents, they'll update an existing registry. We have to go in and just put pointers to the updated documentation.

This is a lot of clerical work, review work. Next slide.

So this is that front page at www.IANA.org/protocols. And there's over 2800 registries that we maintain. That's a lot of registries. We list -- right up here you can see the -- there's an alphabetical listing right here. We've got all the different types of protocols, areas, and then the RFC that defined it as well as what the registration procedure is.

So all the information is there. Take a look. Next slide.

On a day-to-day basis what does actually requesting a protocol parameter look like? You can either send email to IANA at IANA.org and request a protocol parameter. Or we also have a forms page on our Web site where there's different forms for the more frequent types of registrations. So an individual would fill in that form. And what we do is then look at the registration policy for that registry that they're requesting the parameter in. And we find the defining RFC that determines the policy and double-check we have all the correct information.

And then we go into a processing and evaluation step. We follow the process that was set by the IETF for that registry. If required, there is a registration procedure called expert review where the Internet Engineering Steering Group, the IESG of the IETF, will appoint designated experts, technical experts for that particular protocol area. And we consult with them, if needed, have them review requests and actually approve them.

We gather more information from requesters, if needed. And, if, at any time, we have questions about that particular registry, again, we work very closely with the IETF community and the IETF leadership to make sure we have all the information we need to be able to process these requests.

After the evaluation, which in most cases is pretty quick, we update the registry. We have all the registries in XML. And we notify the requester that, hey, your request is complete.

Next slide. Again, this is another view. Media types are something -- those protocols are requested quite often. The request comes in. We

review it. It's a template form on our Web site. We check to make sure everything is complete. If not, we go back and ask them to fill some information in. In this case, the RFC that defines the media type says we have to do this via expert review.

So we will send the request off to the technical expert. They come back to us and either approve or deny or say that we need more information. Once approved, we add it to the registry. We confirm with the requester it's complete. And, yay, they have media type registration they can now go do fun things with.

Next slide.

So how many of these things do we do a month? We do a lot. This chart shows the last 12 months. And this actually doesn't include private enterprise numbers, because it's not in our -- it's in a different statistical program.

So you're looking at a lot of requests that our team processes on a monthly basis. These range from media type support numbers. These are document reviews. These are all different types of things. So we're working hard. Really hard.

Next slide. Performance targets. We worked collaboratively with the IETF to come up with some performance standards. And they actually submit the MoU that exists between ICANN and the IETF. So since 2007 there's been an amendment to the MoU in the form of a service level agreement: In that agreement we have performance standards on our goal times, how long should we be spending on each type of these

requests. And we've been reporting on them since 2007. You can actually find all the reports online going back on our Web site.

Every year the SLA is reviewed. It's modified. We add deliverables depending on new projects that we have that we want to get some additional work done. And they're approved annually.

This is a quick snapshot at looking since January 2014. The SLA met line is the SLA that the IETF sets in the SLA. And we've met it. And, actually, we've consistently met it for many, many, many years. We also have an internal goal that we -- that we set for ourselves. And we've consistently met that. See a lot of 99s, few 100%s in there. Basically, we are doing our work, and we are doing it well. We are doing it timely. We're following the instructions that were given in these RFCs. And the department is busy and doing a good job. And I think I get to pass the mic on to somebody else now.

ELISE GERICH: Before Michelle passes on the mic, are there any questions about the protocol parameters piece? Great. Here.

>> Hi. Maybe this is premature, but I was wondering under what conditions might conflicts arise, if this hasn't been asked already? And how would you handle mediation of those conflicts?

MICHELLE COTTON: Do you mean conflicts if somebody doesn't get their parameter registration or something like that?

(Speaker off microphone).

Okay. So we -- if there is a problem with a protocol parameter registration, we work with the IETF leadership to be able to solve any issues. In general first come, first served registrations are assigned, you know, as they come in, they go out. So I've never seen a problem -- actually, I've never seen a problem really with any protocol parameter registrations. But, if there was a problem, we would seek advice from the IETF leadership. Again, they set the policies that we're just following. If an individual disagreed with a technical expert's review of their request, again, we would work with the IETF community, the leadership, to say the requester disagrees with the technical expert and probably give it another review and work it out that way. Haven't seen any issues in 14 years.

ELISE GERICH: Okay. Michelle keeps mentioning the IETF, IAB leadership. And we have both of those individuals here. We have Russ Housley over here, IAB chair, and Jari Arkko back there, IETF chair. So, if you have any IETF-specific questions, you know where to find them.

Next slide, please, Amanda.

We're going to turn the microphone over to Naela Sarras, and she'll talk about the number resources area.

NAELA SARRAS: Thank you, Elise. My name is Naela Sarras. I've been with ICANN for -- I was reminded this morning -- nine years, almost 10.

And today I'm going to work on -- to talk on the numbers area. So at the IANA department, I work with a team of four people that are all here today. And we process requests for all three areas that will be talked about today.

So we do protocol parameters, as Michelle just explained them. We do number resources. And then we do the domain name requests that come to IANA.

So, Amanda, will you go to the next slide, please.

So this slide is just to demonstrate where our work lies within the unique identifiers. We work specifically on three types of allocations -- IPv4 addresses, IPv6 address, and AS numbers. And that's all our area touches in unique identifiers.

Next slide, please.

Before we go anywhere with explaining what we do to allocate resources, it's really important that we explain that our work here is very deterministic. There's no interpretation of what we need to do.

We allocate IPv4 addresses on a scheduled basis. And we'll talk about that in a little bit. And we allocate IPv6 addresses and AS numbers based on requests from the RIRs.

The data -- by the time a request comes to us, all the data that warrants or justifies the request has already been published. And the decision, if you will, has already been made. This is possible because the RIRs publish their usage data. And then ICANN takes that data in and compiles it and also publishes it on stats.research.icann.org.

So anyone can go up on this Web site and see what the usage level is for the RIRs and see for themselves that they -- a request is warranted or not.

So that's really what the staff's role is to see is it justified? If yes, then it moves on and go ahead and processes the request.

Let's go to the next slide, please, Amanda.

Okay.

So we made this a little bit -- this is scientific.

So we have three different allocation types that were listed here.

The first two are governed by ASO policies. And that's the request that we work on for the numbers area.

The third one is governed by IETF procedures, and that's what Michelle has already spoken to.

So for the first one we're calling formula plus request. And we'll show you the formula. It's, basically, saying how do we go about calculating the space that needs to be allocated? And for IPv6 and for AS numbers, it's based on a request that comes from the RIR.

For IPv4 addresses, it's based on a schedule. That's why we're calling it a formula plus schedule.

Next slide, please. Okay.

So we're going to walk you through what we do for an IPv6 address, for example, for an IPv6 request. Up, up, one more.

Okay. Fine. We won't walk you through that. Okay.

So, when we have a request for an IPv6 allocation, what we check for -- and oh, by the way -- sorry. Before I walk through this, we're giving an example here of what we do for IPv6. The work we do for AS numbers is pretty much the same.

It follows the same structure, the formula or the calculation is, obviously, a little bit different. But the structure is pretty much the same. So I'll walk through the IPv6 address.

So we have our role, as staff that are processing requests, is to confirm that it's coming from an RIR and that they qualify for the space that they're requesting. Up here on the slide it shows you the two methods by which they can qualify.

Once we -- and this we're using again because we have the data all posted on stats.research@icann.org. So that's what we consult to see whether or not they qualify.

Once we determine that they meet the first two boxes, the request comes from an RIR, it qualifies, then we go ahead and we do the allocation.

We told you it's a formula. This formula looking box is just really telling us how much the allocation should be.

And this is lifted straight out of the policy document. So, really, there's not much interpretation here. And that's why we said our job is very deterministic in that matter.

So that's what we do for IPv6 addresses. And, again, pretty much the same thing that we do for AS numbers. Formula is a little different, of course, because the values are different.

The next slide.

Okay. So I mentioned at least two or three times already that we consult stats.research.icann.org. Because, really, that's where we use the data that we use to make our judgment.

This is coming straight out of that Web site. It's an example of IPv6 pool of RIPE NCC IPv6 pool as of October 2nd. We played a little with the colors, but this is coming right out of that Web site. There's other interesting data there. It lists data for all the five RIRs. And it's good to consult. I realize we don't have that link listed. And I think what we can do is update the web page where the meeting session is a related resource. We'll do that after the meeting.

Okay. So we've gotten a request. We've verified that it qualifies for an allocation.

Now what we do next -- please, Amanda, go to the next slide -- is we actually make the allocation based on what the formula told us how much to allocate. And we update the registry.

Here we have a picture of what the registry looks like.

You can see that the -- we walked you through an IPv6 address allocation. Here you see the latest IPv6 allocation that was made by ICANN was all the way back in 2006.

The registry is very -- this registry we played a little bit with the format just to make it easier for the presentation. For example, there's a column on the right missing for notes. But pretty much this is what the registry looks like. It's sortable. It's structured. You can go look at the registry, if you want to.

So we make the allocation. We update the registry. Then next -- next slide, please. We have to communicate the output, firstly, to the requesting RIR, of course, and then to the operations community. And that's our communication piece.

And that's pretty much -- that was -- that's what we do for IPv6 addresses. Again, we do exactly the same thing for AS numbers.

Next slide, please.

Okay. So that was the formula plus request. We receive a request. We verify it. And we make the allocation.

The next type is what we're calling the formula plus schedule. There's the schedule there because this is the assignment of the last of the IPv4 addresses. We do it based on a schedule. Twice per year. It happens on the first business day of March and September each year. So that's the schedule piece of it.

The formula is actually a little software tool that's posted up on GitHub. And we gave you the link. We download it. We run it. We find -- it tells us what allocation to make and we go ahead and make the allocation. So, again, this tool is available. You download it, run it, know the results for yourself before we even get to doing that on the designated day in March and September.

So, again, there's nothing -- there's no judgment that we're applying here. This is all predetermined. We're just processing the requests. And so, again, as we did with IPv6 addresses, we make the allocation on those designated days for the schedule type. And then we publish the results. We communicate them and both again to the RIRs and to the operations community.

So that's for the schedule type of request.

And can you go, please, to the next slide. Okay.

So, in terms of how much -- and like the slide that Michelle showed that showed numbers in the hundreds, we take it easy. We pace ourselves.

So also we really -- it's a very low volume part of our work. Up to now in 2014 we have a handful of requests. And I don't see reasons for having many more this year, right? So maybe it will go up by -- I don't know. But, essentially, this -- it's a really low volume part of our work.

So that's our allocations per year.

Next I think we talk about performance standards like Michelle did.

So we had a consultation with a community with 2012 about what performance standards to track against and how to set our key performance indicators.

Since the beginning of tracking our data against those KPIs, we've been meeting or exceeding our targets. On the slide -- I don't know how well you can read it -- it lists what our key performance indicators are. And then we have a whole page on the IANA web site dedicated to what our KPIs are and the actual data of how we're actually doing against the

KPIs. So how we met the KPIs in specific requests that we processed. And I think Elise will show you that page at the end of the presentation.

But, yeah, that's what we're doing for performance targets. We track ourselves. And I think that's pretty much what we wanted to speak about in the numbers area. Thank you.

ELISE GERICH: Thanks, Naela. Are there any questions for Naela about the numbers area? Any remote participation questions? No? I do want to mention that, when Naela said, you know, we have low volume, you've got to realize we had really large blocks. So we've done one allocation of IPv6. And no one's run out of them yet. So they haven't come back for more. Typically, our requests are -- have to do with autonomous system numbers. I did receive a handwritten note question.

And that was "Would it be possible to describe what the acronyms mean before we just talk in letters?" So, if you'd raise your hand if there's an acronym you'd like us to answer right now, we could try and put that out of the way. Anybody -- did you know what AS number meant? Autonomous system? Okay. IPv6, means IPv6 addresses. All righty. Well, we'll move on to the next -- go ahead.

NAELA SARRAS: Perhaps you can explain how we structured this presentation to be about what we do, not the theory of what we do, maybe, a little bit.

ELISE GERICH: First of all, Mary.

MARY UDUMA: Yeah. My name is Mary Uduma from Nigeria. I wanted to know if that's a particular template form that the RIR do fill or just send an email to request a block, please.

NAELA SARRAS: Actually, I'm going to ask my colleague, Leo, to answer that question. Leo.

LEO VEGODA: So for the RIRs, there's a template that they can fill out which is published on www.IANA.org. And they can send to us. And it's just so that they make sure they provide all the information that's required. If they sent a video of interpretive dance that also gave that information, that would be fine. But it would take a little bit longer.

MARY UDUMA: Thank you for that response. It would have been helpful if you had shown the -- since some of us have not seeing (indiscernible) is not making we've not been adventurous enough to go there. But would have been helpful if you show what the template is like. And we may want to ask questions why you ask certain questions on those templates. Thank you.

ELISE GERICH: Okay. That's good input for us, Mary. Thanks. And we'll try to improve that the next time we do it.

If there are no more questions, I'm going to hand it over to Kim Davies, and he's going to talk to you about the root zone and the names function.

KIM DAVIES: Thanks, Elise. My name is Kim Davies. I'm director of technical services, but within IANA the primary reason I'm here talking about this is I'm the subject matter expert for domain names.

Next slide, please.

As Naela presented earlier, IANA maintains thousands of unique identifier registries, the vast majority of which are IETF unique identifiers.

When we talk about the domain name stuff we do in IANA, we're really referring to our management of delegations within the domain name space.

There are other domain name related registries we maintain within IANA. For example, resource record types, security algorithms, header flags. There are several dozen more. They are all maintained using the IETF processes. So someone issues an RFC or requests a reservation, and then we update the registries that way.

But when it comes to the domain name space itself, domain names, we have a fairly unique role, which is managing the root of the domain name space.

So next slide, please.

I suspect this is not news to anyone in this room. The domain name system is a hierarchy. Each level of the domain name system is issued to another party by a process of delegation.

At the top of the domain name space is what we call the root zone. The root zone is responsible for being the authoritative record of which top-level domains exist. So, you know, within ICANN this week, we'll talk about gTLDs, we'll talk about ccTLDs for country codes. All of these are different flavors of top-level domains. All of them need to be listed in the root zone for them to become active. And managing the content of the root zone is one of the IANA functions that we perform.

Next slide.

So when we talk about managing the root zone and managing that list of top-level domains, what is it that we actually do?

Well, firstly, we wait to receive a request. We receive requests from top-level domain managers, companies that run the registries for ccTLDs and gTLDs.

The kinds of events that happen that trigger one of these managers to reach out to us, sometimes teams the wholesale change of who runs a TLD. Sometimes it's simply routine maintenance. The designated point of contact from a particular registry perhaps has left the company. They're appointing someone different to replace them. We need to know who the correct representative of that organization is, so they need to update our database.

There's technical changes. So if new servers come online to operate that TLD, those new servers need to be listed in the root zone.

There's also, and it's relatively rare but it does happen, in the event that there's some kind of natural disaster, perhaps there's a hurricane and earthquake, often what we see is that it severely impacts the ability of critical Internet infrastructure in a particular country to operate. It might mean, for example, that a large number of name servers within a particular company are affected and are no longer available.

So in the event that there is some kind of catastrophe of that nature, we often are in direct dialogue for the TLD manager for that particular country. They'll reach out to us and have an emergency plan on bringing up new servers, perhaps in new locations, perhaps in different geographical areas, and we work to process that and expedite that.

So next slide.

So what's in this database? Well, it's actually not a very large database. We maintain a few key pieces of information for every top-level domain.

Firstly, we maintain the identity of the recognized operator of the domain. That is, the name, the correct legal name of that organization, its place of business. We maintain at least two contacts for each domain, an administrative contact and a technical contact. For each of those contacts, we have their name, their job title, address, phone, fax, email, everything we need to contact them and some archaic ways of contacting them as well.

We maintain the technical configuration for the top-level domain. That involves two key pieces of information. One is the list of name servers for a particular top-level domain, and the second is the DNSSEC records.

So when you sign a domain name with DNSSEC, it produces a string of letters and numbers. That string of letters and numbers that we call a DS record needs to be inserted in the root zone. We store that as well.

Finally, we maintain some meta data for every top-level domain. This meta data is basically informational. We list it as a courtesy. It's something that's useful for the community, but we don't actually need it to do our jobs. So that includes things like what's the Web address of the registry for a particular TLD? Where is the WHOIS server for a particular TLD? Things lake what's the translation of an IDN into English? That kind of stuff.

So we have a few records there. We're looking to expand it, in fact, in the near future to add things like translations, and so on. But it's essentially additional information that's useful to the community, but it's not core to the operations of the TLD.

So next slide.

So here's an actual example. Rather than talking about hypotheticals, I went on our database a few weeks ago and pulled this record for .HAMBURG. Here you can see it has the legal name of the entity that runs .HAMBURG. It has the address. It has the two contact persons, their contact details. It has their technical configuration details. Tells you that the Web site for .HAMBURG is at www.dothamburg.de. And their WHOIS server is whois.nic.hamburg.

That's really it. And all that information is public. You can go onto the IANA Web site and see it in a section called the Root Zone Database.

Alternatively, you can actually go to WHOIS, whois.iana.org, and do a query for a TLD and find that information.

So nothing there is secret. That's all public information.

Next slide.

So we get this request or there's a need to generate a request for the reasons I described earlier. Generally speaking, our customers lodge it by an automatic system. We invested a lot of time and effort a few years back developing a system we call RZMS, the root zone management system. Essentially it provides a Web portal where all TLD managers are issued with a user name and password. They can log in, they can lodge new requests, they can check the status of existing requests. It provides basically self-service for them to do most kinds of change requests.

Next slide.

So a TLD manager has a need to make a change. They've submitted a request. What happens next?

This is really the bulk of where we, as IANA staff, do our work when it pertains to these kinds of requests.

Our job is to look at that request, make sure it meets with certain policy and technical requirements, confirm that the parties that need to consent to the change agree. If we find issues during that processing, we clarify with the applicant, the requester, to find out is this an issue that can be solved. If not, we'll close the request. If so, we give them some homework. They go off and fix the issue and then we continue.

Once we're satisfied that all the policy and technical requirements are met, we then forward the request to NTIA to authorize, which I'll get to in a moment.

So next slide.

So what are the things we're actually looking for? Firstly, technical. And this is key. We don't want to insert entries into the root zone that will break a TLD, that will cause harm to other TLDs because it somehow affects the function of the root zone.

So there's a set of technical requirements that were developed in consultation with the community a number of years ago that form the basis of the checks we do on a technical level. These include checking that all the name servers for a TLD are actually working, do they respond, can we send queries to them and get back answers that actually make sense.

When they sign their zones with DNSSEC, we're check to go see can we actually do a proper validation using DNSSEC. DNSSEC involves, we saw on a previous slide, a string of letters and numbers that are perhaps not really intuitive to see does that look right. The real way to test if someone made a typo in that, is to do a validation, do all the computation involved, make sure it all lines up.

Email addresses are an important part of our work flow. A system sends our email to us for verification, and so on. So whenever a new email address is added to our database, we want to make sure it works so there's a process of e-mailing them, validating their email address is functional. They get emails from us; we get emails from them.

The next segment of the review is consent. Generally speaking, we want to guard against the wrong party getting some kind of control of a TLD, so this process involves essentially getting all the existing people that responsible for the TLD to agree to any future changes.

So when we add a new contact to the domain, we make sure the existing contacts agree. If a new contact is coming on board that hasn't performed that role, we make sure they understand that they're accepting of this new responsibility.

There are certain types of technical changes that we do where the name server that's being modified actually impacts multiple TLDs. So this is more common in the past, but some individual servers actually supported 40, 50, 60 TLDs on them.

We want to make sure that if there's a configuration change to that one server, that no one is surprised. So we have a process of ensuring agreement from other TLD managers when we get those kind of requests.

We have a regulatory check. We have certain legal requirements to follow the law. So we make sure that anything we're doing in the request complies with legal requirements.

We check for well formess. We don't want to have a database that's full of junk records. We want to make sure they're accurate, they're clear. Where it makes sense, we want them to be in Latin script so people can read them. Some people submit their entries in all caps. I think it's a little more pleasant on the eye that we make it all a bit normalized. Really there we're just making sure it's all consistent and tidy, so when

you look at the whole database it looks fairly clear. And I'll get to more detail I think on the next slide. But, transfer of responsibility.

Whenever there is a significant change in who operates a TLD, a number of additional policy requirements kick in.

Essentially, as long as, you know, these early steps are complete, any kind of routine maintenance update to a TLD just needs that, just needs the existing contacts to agree, needs to check out technically, and more or less that's all that's required. It's a very quick process. You know, these changes we do day in, day out, very simple.

But if you're transferring operation of a TLD from one whole entity to another whole entity, we need to do additional processing as well.

So next slide.

So the process of doing a transfer of a TLD, we call this redelegation, differs whether it's a gTLD or a ccTLD, and they differ significantly.

In the context of gTLDs, really the ultimate driver of this process is GNSO policy.

GNSO policy at an operational level is affected by the contracts that registry operators sign with ICANN.

So within IANA, what we're doing is basically looking to make sure whenever there's a transfer to a new TLD operator for a gTLD, that that is the current contracted party, the most recent contracted entity with ICANN is correctly reflected in the IANA database.

So the key point there is we're not making a judgment call on who should run it, under what terms. That happens outside of the IANA department, elsewhere in ICANN. There's a whole gTLD department responsible for that.

What we're simply looking for is when that process is all said and done, when there's a freshly inked contract with a new party, that we reflect that in our database.

Now, if we look on the right, ccTLDs. Now, ccTLDs perhaps counter-intuitively, we don't make decisions there either. We have a process of review, but what that review is essentially based on is was there a consensus within the country to make a transfer. And we have a number of evaluation criteria pertaining to that. We ask for evidence that the consultation has taken place. We ask for statements from various parties within the country, such as governments, private sector, academia, and so on. But really what we're looking for there is there was a process conducted within the country to determine who should operate the ccTLD, and that the change, transfer, or redelegation happens, reflects that concensus that happened.

So next slide.

So once we've done all that processing, we're confident that a request should proceed, we transmit our recommendation to NTIA, the U.S. government role that is responsible for the IANA contract. They review what we've provided and they provide authorization to proceed.

Based on that authorization, implementation happens. So this involves multiple parties. Firstly, the root zone file itself is maintained by

VeriSign. So the authorization for any changes, those technical pieces gets transmitted to VeriSign as well. VeriSign implements the change based on an authorization received.

But that WHOIS change, like if a contact changes or some of those other elements changes, that's something we will implement, so that authorization will allow us to act to update those other records as well.

Once the update is made, you can see it online. Again, just like everything we've been discussing today, it's all public information. There's no secret database. It's all online.

So next slide.

So I'm going to switch gears slightly.

Another critical piece that we manage in the context of domain names is the root key signing key. This is something that we took on starting in the year 2010 with the introduction of DNSSEC for the root zone.

I don't want to go into a class on how DNSSEC works, but suffice it to say there's a key file. We call it a key but it's a computer file that contains codes that are used to generate the signatures that ultimately are used to perform all the verifications involved in DNSSEC.

The way we do it is we follow a process of performing key-signing ceremonies. These ceremonies again are public. They're designed to be auditable. They're designed that the community trusts that they've been done in a proper way. And let me explain a little bit about the detail of how those work.

So next slide.

So the heart of managing the key-signing key is a little box. The box is a little smaller than a laptop, a little thicker. It's called a hardware security module or an HSM.

This box is designed to store a file very securely. In fact, it's a really interesting thing. It's like a hard drive, but it's designed that if you try and open it, it self-destructs. If the temperature goes too far off the chart, it self-destructs. If you shake it too much, it self-destructs. It's really designed that any attempt to sort of get the information out of it in an unauthorized way will result in its own demise.

So we store the key signing key in these devices. We have multiple for redundancy but each one contains the same information.

Next slide.

So if you try and tamper with it or try to get into it, it self-destructs. Well, how do you use it when you want to use it?

Well, it uses a series of smart cards. You insert smart cards in, and once enough smart cards have been inserted, it comes online. It's activated.

So the way we've configured these is there are seven smart cards to operate each one of these, and you need three of those seven smart cards to be present in order to switch it on.

Next slide.

Now, this is key. In order for this system to be secure, the way it's been designed is that we don't have the ability to turn these things on. What we've done is we've taken those seven smart cards and we've given each one to a different community member. So in order to use one of

these HSMs, we need to get, from this pool of seven community members, at least three to come to a ceremony, be present, provide access using the smart card that was issued to them. This will switch on the HSM and allow us to do the key signing work that we need to do.

So next slide.

So where do we store this HSM? Well, the first thing we do is we put it inside a high-security safe. The combination is only known to two individuals. We call them safe security controllers. The safe security controller knows nothing else.

They're not involved in operations other than they know the combination to the safe. So we store it in a safe. That safe is monitored with a seismic sensor, other kinds of sensors.

So next slide.

So we take the safe -- there's actually multiple safes for redundancy and also we need to store some additional equipment there as well. We store it in a high-security room. To get access to this room, there's two more people that need to be involved. One is called a ceremony add administrator, the other is called an internal witness. There's a bio metric scanner on the door so for us to get into the room, these two individuals need to be present as well.

So so far we have at least three community members, there needs to be the person who knows the combination to the safe, two more people that need access to the room.

So next slide.

This room is in turn stored in a larger room which we call our ceremony room. Inside this room is where we actually do the work involving those HSMs so we pull them out of the safe and we bring them out into this room.

When we're in this room, everything is on camera. Everything is auditable. We have a Web stream so you can see what's happening in this room when we do this.

We have a bunch of witnesses. There's usually at least 20 people there present when we do this. And access to this room is controlled by someone else that is not present at all.

So, in order for all the people I just talked about to be present, someone else that's off-site, that's remote, needs to call into the security company, disable all the alarms to give a list of all the names of all the people and their full names on their identity cards to the security guards at the facility. So that, when we come in to do a ceremony, they're checked off a list. Everything's deactivated and ready to go.

So the reason I'm telling you this is that, you know, there's an orchestrated process of conducting these ceremonies and getting access to that secret file needs to involve many different people. These are part of the main security considerations that are involved in managing the key signing key.

So next slide.

So we have two facilities, more or less identical to each other. One on the West Coast, about 30 kilometers are from where we are today in El Segundo, California. In fact, it's just south of LAX airport. We also have

one in rural Virginia over towards the east coast of the U.S. in Culpepper. Both of these, basically, can do the same job. They're there for disaster recovery, in essence. If we had one site and there was a problem, then what would we do then?

What we do in practice is we alternate. So we do four ceremonies a year. We alternate East Coast, West Coast, East Coast, West Coast.

And I mentioned before that these facilities are guarded. They're secure facilities by a third party. So no one can just walk off the street and get into these facilities. Next slide.

So I mentioned four times a year we get all these folks, these trusted community representatives, the staff in those various roles. We come together. We pull those HSMs out. We connect them to a laptop. And, in essence, what we do is we generate three months' worth of security signatures. And those three months' worth of security signatures reduced day-to-day by VeriSign in issuing the root zone file itself.

I mentioned the process is streamed. It's recorded. There's literally hundreds of hours of video on our Web site, if you want to amuse yourself with looking at the last 19 ceremonies we've conducted. All the videos are public. The code that we use on the laptop is public. The scripts that we follow to enter the facility is all there in pdf files. So really everything that's involved in this process is transparent. And that's important because the whole reason we do this is to trust the process.

The way security works is that, if there's a compromise of the private security data, then it's not really security at all. So the community as a

whole needs to trust that the files inside those devices hasn't been tampered with or has been accessed by anyone than those that were expected to use it. So the TCRs that I mentioned, the trusted community representatives, are from the community that's here this week. But they're definitely security specialists. They're people that know when we've done something wrong. They can monitor the situation. Basically, they're there to provide onsite advice and make sure that we do everything by the book, that we don't do something in the wrong way. And then they can go back to their communities afterwards and say I was there. I saw it. It was done the way it should have been done. That device has not been tampered with.

Next slide.

So this is the attendance at the last ceremony that we conducted. And I've just walked you through the -- sort of the high level overview of how it's constructed. But it's probably -- you're going to get a better sense of it if you look at at least one of these two documentaries. The last two ceremonies we did, we have documentary film crews there. The first one was from The Guardian newspaper. They produced a story as well as some video footage on their Web site. More recently we had the BBC there. They filmed a documentary called Horizon. There's about -- both of them I think are 5 to 10 minutes long, the video footage. And I think you'll get a better sense of how those ceremonies are conducted if you wanted to go look at those videos at some point in the future.

So next slide.

So, consistent with my colleagues, I'm finishing up the sort of section on domain names with some of our performance reporting. We produce a lot of performance reports. They're on our Web site, IANA.org/performance. This is just a summary.

It's just, again, to give you a flavor like the other ones that we measure how long we take; we measure how many requests we receive; we produce monthly reports. We have SLAs, those service level agreements, those targets agreed with the community-based on a consultation we did two years ago. And, as you can see here, we do typically somewhere between 50 to 100, more or less, requests per month in terms of updating the root zone database. We have -- you know -- we have percentages in terms of KPIs that are met. I think Michelle and Naela get it easy in their sections because the great thing about what they're doing is it's highly dependent on IANA staff doing their job in a timely fashion.

One of the more unique things about managing the domain name space and particularly those transfers I referred to earlier, a lot of the work that goes into getting the documentation together for transferring a ccTLD, for example, involves the requester going back into their country and doing work. So a lot of the timing in that is not entirely predictable. So we have a lot of back and forth. Sometimes it goes over many months. So some of the, I think, the percentages are a little skewed by that. But that's a normal operation. But importantly, we report directly to the ccTLD managers, the gTLD managers about our performance. And, you know, as we move forward, we continue to look and revise our SLAs.

So next slide. I think that's it. So thank you.

ELISE GERICH: Does anyone have any questions for Kim on the domain name portion of our presentation? Okay.

So one. Thank you.

SERGE RADOVCIC: Serge Radovcic, RIPE NCC. Just for a bit of clarification, can you go back to the last slide for me.

KIM DAVIES: Sure. Amanda.

SERGE RADOVCIC: With the performance review, when you talk about the end-to-end days, you said that the NTIA only gets involved when there's redelegations? Or they also involved in these processes?

KIM DAVIES: They're involved in all these requests. Every change in the root zone database needs NTIA authorization.

SERGE RADOVCIC: What would you estimate their involvement is in that -- in that time span of the end to end?

KIM DAVIES: I think it's relatively minor. You know, we usually get a rapid turnaround from NTIA. It's not a significant portion of that time period, I would say. You know, I don't have the numbers, averages off the top of my head. But the bulk of that time is typically requester time.

Generally, with a few rare exceptions, you know, the ICANN piece and the NTIA piece and the VeriSign piece together is, you know, two, three days of that for a normal request max.

The rest of the time is often going back to the requester. There's like a technical check problem. And they go and remedy their configuration. They try again. They take some days, so on. So there's a mix of those things.

MATTHEW SHEARS: Yes, Matthew Shears, CDT. Just a followup on the gentleman's question over there. Do you get questions back from NTIA, requests for additional information, any kind of queries, any kind of comments back about the process?

And the other one is you seemed to imply ccTLD redelegation, that you are making a judgment of sorts. Could you explain that a little bit more? Thanks.

KIM DAVIES: NTIA comes back with questions from time to time. I would say I've been at ICANN for nine years now. And it's rarer now. Because I think the industry as a whole has matured. ccTLDs and gTLDs are familiar

with what they need to provide. We're familiar with the process. It's just better down.

So it's really -- everything is very quite routine. So it's extremely rare in the modern context we would get queries from NTIA. Because we pre-emptively know what they would like to see when we submit a request for reauthorization.

In terms of ccTLDs delegations, we're not trying to make a judgment call. We're trying to, basically, make sure that there's consent to make a change. You know, this is one of the fuzziest areas that we work in. Because we're working basically off principles that were defined in a document written by Jon Postel 20-odd years ago. But the communities recognize the challenge we face here. So there's been an ongoing multi-year effort with the joint ccNSO GAC working group called the framework of interpretation. I think meeting, in fact, they're planning to come up with their final ratified recommendations. But, in essence, they're trying to provide objective definition of what those principles are and how IANA should assess it. So we're hopeful, as a result of that, we're going to be revising our guidance on how ccTLD redelegations or transfers happen. And that will provide a much more objective process effectively checklisting it to a large extent. And whatever sort of judgment we might have needed to apply, you know, is decreasing. And we're trying to get it as close to zero as possible.

But it's not binary. So there's always going to be on ongoing evolutionary process of looking how we do it, consulting with the community and making sure that, you know, corner cases are identified, that there's a proper process that kicks off that the community knows

about and new issues that are identified. And they provide guidance to us on how we should deal with those kind of issues.

PAUL FOODY: Hello. Paul Foody, speaking on my own behalf.

I apologize if you've answered this already, but the sound down at this end is pretty difficult, particularly with the translators behind us.

What I'm asking about is the DNSSEC. Did you say there were 19 that have been signed so far?

KIM DAVIES: There's been -- I think maybe 18. 18 or 19 ceremonies. So we do a ceremony every three months. We started doing them in June 2010. So, at a rate of four per year, I think we're up to the 19th one.

PAUL FOODY: And is that a signing ceremony for each registry, each TLD?

KIM DAVIES: No. So what happens is we generate a set of signatures that are used that have a validity of a couple weeks each. We generate three months' worth of signatures that VeriSign uses to apply to their daily root zone updates. So it's -- you know, it's highly technical. But, in essence, it's a set of files that VeriSign can use to produce daily root zone updates. But they're not tied to the TLDs or the contents of the root zone itself. The contents of the root zone is an entirely different process. This is

really about the signature that wraps around the whole root zone on its daily issuance. Does that make it clear?

PAUL FOODY: I think so, yeah. How will that be affected by the massive increase in the number of TLDs?

KIM DAVIES: It's actually not affected at all. It's completely level.

PAUL FOODY: Thank you.

KIM DAVIES: Please.

MARY UDUMA: Thank you. This is Mary again from Nigeria. And I just want to follow up on the question on ccTLD redelegation and the updates. Looking at this slide as well, is it possible to get a sense of what -- how it -- how long it takes to process a ccTLD, a ccTLD in delegation or updates as against the order names, you know? When you have registered, it shows this. But, when it comes to cc in particular, it might be a different graph.

And do you have SLA for the redelegation only related to ccTLD? Thank you.

KIM DAVIES: We do.

So that slide was, you know, a one-page summary of a raft of measures that we maintain. And, in fact, we publish very long detailed SLA reports on a monthly basis. So the precise things you're asking for we do report separately. It's in our performance reporting at IANA.org/performance. But Elise has kindly flipped forward to one of her slides, her summation slides, to give you a taste of it. So here is root zone processing broken down into type of change requests. So we had, for example, in the reporting period here, 19 changes to the name server records, 2 changes to DS records, 12 changes to the admin contacts, 9 to the tech contact and so on. 34 delegations and redelegations for that month. And then we have the number of days on the right.

Now, this is averaging gTLDs and ccTLDs together. But, in our even more detailed reporting, it does break it out even further.

And we have different SLAs based on how long we should typically take for ccTLD redelegations and delegations versus gTLD delegations and redelegations.

So we definitely have that information. And, certainly, if that information is not comprehensive enough, then please talk to us and we can walk you through it and see where there might be gaps that would help you. Any other questions? Please.

BEN TOUNG: Hi. My name is Ben Toung (phonetic) and speaking on my own behalf. Trying to get a better understanding of the government oversight

function here and how much of an active role do they play in that approval? Are they functionally acting as a rubber stamp, or in what cases have they denied the requests? Or what type of questions are they typically asking for what information on which to base their decision? Just trying to get a better picture of what that oversight function is in practice.

KIM DAVIES: I mean, I don't want to speak for them. I don't want to put Ashley on the spot right now either, unless she wants to.

But I think -- you know, it's relatively rare. And, you know, compounding on what I said a few minutes ago, I think, particularly in the last few years there's a much better understanding of what they're looking for.

And I point to the new gTLD delegation process as a good example. We had constant dialogue with NTIA prior to delegating the first new gTLD. We wanted to understand what they're expecting to see when it came to delegating new gTLDs. We produce a report now that specifically responds to all the items that NTIA wanted to see in terms of authorizing a new gTLD delegation request. And we have hundreds of new gTLD delegation reports on our Web site. I invite you to go have a look at one. But it was really tailored to what's the precise amount of information they need to see to authorize a new gTLD being added to the root zone.

So that's kind of the direction we're going in. And, you know, I think -- you know, again, I would say NTIA requests for clarification are

extremely rare now. And that really talks to the maturity of the process. There's really no surprises in the root zone management process any more. ICANN has been doing it now for 15 years. We think we know more or less all the modalities that are going to happen. We have a well-understood process. The community is quite familiar with it. At least the direct TLD management community that makes use of the process.

And, you know, I think that reflects in the fact there's really very rarely any hiccups in the process going through.

Thank you.

GAELLE FALL: Thank you, good morning. My name is Gaelle Fall. I am here with -- I'm here from AfriNIC and Mauritius. And my question is in regards to number resources, if I can maybe direct my question to Naela or whoever can answer.

Seeing that the IP version 4 is -- space is being exhausted and that we are now going to be moving into a more abundant resource, how will the number distribution mandate change for IANA?

NAELA SARRAS: Thank you. So I might have to lean on my colleagues here, because we're really just processing requests.

So I'm not really sure I understand that question. There's a policy upon which we -- we allocate resources, and any changes for the IANA mandate would have to come to us as a policy change. So right now,

we're under the last policy that we're working under, and that 2006 allocation you saw is what we're working under.

And the stats Web site that we showed you shows you who has what resources still remaining and when they would qualify to come ask for more.

So I'm not sure if I addressed your question, if you have something more specific that you wanted to ask.

GAELLE FALL: Well, specifically, the issue is that you will come to a point where seeing that your -- the IPv6 that you've allocated -- you've allocated them since 2006. We are now in 2014. Once the IPv4 is exhausted, IPv6 is going to be the only resource that's going to be allocated.

So if you're going to be allocating IPv6 blocks once every ten years, I would assume that it's -- I mean, this just -- I'm just saying.

NAELA SARRAS: Oh.

Because it is an abundant resource.

So I'm just wondering whether it's going to change the actual -- the actual number part of the -- of your function, of IANA's function. I would assume it will change it in some kind of way seeing that you will not be allocating as many blocks as you would have done with the version 4 resource.

That's basically my question.

ELISE GERICH: Yes, you've hit the nail on the head. And as you could see from the slide deck that we put up earlier, we've only allocated five number resource blocks in this last year. And one of them was an automatic formula for the IPv4 returned pool.

Right now, the IANA functions department has no IPv4 blocks to assign. Everything is -- that we have in our IPv4 pool has been returned to a Regional Internet Registry who have then returned it to the IANA functions department. And based on the global policy from the five Regional Internet Registries, we have a formula. And on September 1st and March 1st we apply the formula, and however big the pool is, that formula runs against the pool and they get allocated.

So that's all we do for IPv4 now. So at the max, there's two allocations. And there's hardly any humans involved in that at all because it's a formula and it's a program and it happens.

IPv6, as you said, are none of those five allocations that we made this year. They're all blocks of autonomous system numbers.

So I don't think we're going to see much of a change, even if there's nothing in our recovered pool, because autonomous system numbers are what are really driving the numbers area today.

And as you get more organizations that build their own networks and need an Autonomous System Number for routing protocols, for

aggregating their networks to talk to another network, that's where their numbers are.

But we track that very closely. And I think as we mentioned earlier today, Naela runs a team of four people, and those four IANA specialists are responsible for all the requests across all types of unique identifiers. That includes protocol parameters, root zone changes as well as the numbers.

So I think numbers are probably going to stay flat for quite a long time.

Were there more questions?

Yes, please push the microphone.

VALENTINA PAVEL BURLOIU: My name is. Valentina -- Is it working?

My name is Valentina. I am from Romania. I have a legal background, so I'm not technically skilled at all, but I read the SSAC report and I found out that there are actually 11 IANA functions, but only four of them are in the contract with NTIA. And I was wondering why and what happens to the other functions which are not in the contract.

ELISE GERICH: Okay. So the SSAC document is an excellent document and talks about all the kinds of things that a department within an organization might do.

So the IANA functions department within ICANN does some things for people outside of NTIA.

So, for instance, with the IAB and the IETF, we have an agreement to manage the time zone database. That's not considered IANA functions within the contract that we have with the Department of Commerce.

So we're a department that can do many projects, and ICANN has the contract with NTIA, and so our department is the one that also delivers on that contract as well as some other activities.

And the IANA functions contract does have some other services. One of them is the .INT domain. So we act as the manager for the .INT domain. But that is not something we've talked about today because it's a very small part of what we do.

Does that answer your question?

Okay. Great.

Do we have any more questions from the audience?

We really appreciate your engagement by asking us this.

So we've had our summary of the protocol parameters, number resources, and the domain names. So if we could go to the next slide, please.

This goes back to one of our questions. How big is our job and what will the numbers mean to it?

So what we have here are some bubble charts. And so basically, if you look at the big bubbles through the middle, that will tell you what we do for domain-related requests, number-related requests, protocol-related

requests, and then general questions we get. Much like the last question; you know, what else is part of the IANA functions?

So those are the numbers of types of requests we get.

Then we have some other things. We have two third-party audits that we run every year. One of the audits is for the DNSSEC security of the key signing key. Another audit that we run is for the IANA systems in general. We have a contract with Price Waterhouse Cooper and they run these third-party audits for us. We publish the certificates we've gotten for the DNSSEC audits and the SOC2 audits that we get for the SysTrust is an internal document so that we can improve our own systems.

And then we notice the four key signing ceremonies. Four means -- sounds like a small number, but if you listen to Kim's presentation, you realize that a lot of people are involved. It takes a lot of coordination. We have two cryptographic key managers that have joined our department. They are both here, Andres and Punkie (phonetic), so they are in charge of making sure that those key signing ceremonies run smoothly and that we keep the key signing key secure.

Next slide, please.

Every year we run a customer survey. Some of you may have received that from a -- How do you pronounce that? Ebiquity? Is that the company, Leo? Ebiquity.

And hopefully you've replied to our customer survey. This is just a summary of the 2013 survey. The 2014 survey is coming to a close this week, and we'll be publishing the results. We obviously look at the

results and we have a continuous improvement manager within our department to help us identify areas where we can make improvements and try to improve the service that we offer to all of you.

Next slide, please.

This slide we showed earlier after one of the questions. This is just on our performance Web page. IANA.org/performance. And monthly, we publish this snapshot that shows you the average processing times for different types of requests on the root zone database.

Next slide, please.

This one is something that sometimes we find that a lot of people have a misimpression of what we do and what we don't do. So we've tried to put it right here in green checks and red Xs.

The green checks are the things we do. The red Xs are the things we don't do.

So I just wanted to make sure it's quite clear that we do create registries, and they're all based on policies that come out of different organizations, different communities.

We maintain those registries. We allocate number resources. Infrequently, but we do do it. And republish all registries for general public use. So all the registries are public. They're available to everyone. There's no charge to use any of these authoritative registries for protocol parameters, numbers, or root zone.

The things we don't do. We don't create policy. We're obviously part of ICANN, but we're an implementation and operational department, and

that's what we do. So maybe it's a little less interesting than you thought.

We don't determine who gets to run a domain. That's done by all different organizations outside of our department of 13 people.

And we don't get to choose who the TLD managers are. It might be a heady example to try to figure that out, but I don't think that's something we'd ever want to do, and we don't choose who runs the TLDs.

Next slide, please.

So finally, just in summary, we maintain unique registries. These are straightforward, and they're not generally known to most end users. I mean, most of our questions today were about the domain names. We had very few on protocol parameters. We had a couple on numbers. But the largest number of registries that we maintain are for the protocol parameters, and very few people are aware of them unless you're a developer or an equipment vendor or someone who's building a network.

Next slide, please.

And I want to thank you for your time and attention. And I'd like those of us in the IANA Department to please raise your hand again so that if someone sees you in the halls, they can stop and ask you a question.

And we're all here available if you have any questions. We've put some Web site information, and we're here to answer more questions, if you have any more questions.

If not, I want to thank you very much for your time, and I appreciate the fact that you're all here in L.A. at the ICANN meeting.

Thank you.

[Applause]

[END OF TRANSCRIPT]