SINGAPORE - IANA Department – Who, What, Why?
Wednesday, February 11, 2015 – 08:15 to 09:30
ICANN – Singapore, Singapore

ELISE GERICH: I'm ready? I'm ready. So it's 8:17 by the little clock that's right beneath me, and I want to welcome you all to the IANA Who, What, and Why session, subtitled, "Why the IANA functions are less interesting than you thought," and I can see from many of the faces in the room that I know some of you but not all of you.

So my name is Elise Gerich and I'm the vice president of IANA and technical services at ICANN, and my colleague, Kim Davies, is the director of technical services for the IANA department at ICANN. And what we'd like to do is just kind of walk through what the IANA department, the IANA specialists who process requests for names, numbers, and protocol parameters, actually do. What our job is.

It seems like we're the topic of a lot of conversation recently, so I thought it would be interesting to make sure that people understand what we do on a day-to-day basis.

So the folks who work at ICANN in the IANA department are the ones you see pictured above. There are 12 of us, and as you can see on the top row, Kim Davies and myself are there. Naela Sarras, she manages the IANA specialist team that process each of the requests.

Michelle Cotton is our engagement liaison with the IETF and the protocol parameters. Michelle just celebrated her 15th anniversary at

ICANN. She was the second employee for ICANN, so she's well-versed in all the aspects of what ICANN does and the IETF would never let her go.

[Laughter]

ELISE GERICH: She's been a wonderful, wonderful ambassador from ICANN to the IETF.

Pearl, Amanda, Selina, and Paula, are the four IANA specialists that process all the requests that come in. They handle the protocol parameter requests, the numbers requests from the regional Internet registries, as well as the root zone change requests.

Andres and Punky have joined us just recently. They're our cryptographic key managers. They work very closely with Kim on the DNSSEC for the root, so they take care of all the key signing ceremonies and things like that, and you'll learn more about that as we get further into the presentation.

And Marilia has just joined us recently. She's our manager of business excellence. We have had an EFQM quality management program for the last five years, and Marilia is taking on that responsibility for continuous improvement within our department.

So just a little history.

And basically I think everyone's heard all week that ICANN was established in 1998 to be the steward and operator of the IANA functions, and basically we have a contract with the Department of Commerce, and our responsibility is to maintain the registries, the lists

of the unique identifiers for the Internet, and those unique identifiers have been defined by the IETF.

And so we collaborate and cooperate and coordinate with the different communities that -- in the distributed system, such as the regional Internet registries, the IETF, the ccNSO, and the GNSO. And these are authoritative registries that are used by vendors, service providers, business applications developers, and others.

Without the maintenance of an authoritative registry, there would be less interoperability, so if I were to be designing a new piece of hardware that I wanted to connect to the Internet, I'd need to look at the protocol parameters registries.

If I wanted to be able to have my router communicate to another router, I would have to make sure that I worked through the regional Internet registries and got allocated an address that could communicate and would be unique.

So those are -- that's the job we have is to maintain these registries.

So the three primary functions that we have, as I mentioned earlier, are the protocol parameters and the registries that support those, the number resources, and the domain names, which is really the root zone change requests.

We're going to talk about those one by one.

For the protocol parameters, I'm going to pass along the microphone to my colleague, Kim Davies, and he will introduce you to what we do with them.

KIM DAVIES: Good morning, everyone.

As Elise mentioned, we tend to divide what we do into three separate areas: Protocol parameters, number resources, and domain names.

But really, at its heart, everything IANA does is about protocol parameters. The domain names and number resources, as you'll find out later in the presentation, are just highly specialized forms of protocol parameters.

Now, underlying the Internet are really hundreds and thousands of different protocol parameters.

They're hidden identifiers that you wouldn't perhaps know are essential to the functioning of the Internet.

Generally speaking, the vast majority of these are only known to software implementers. As an end user of products, you wouldn't necessarily know that they're there. They're not part of the user experience. But nonetheless, coordinating these protocol parameters is essential, as Elise just mentioned, to interoperability of the Internet.

So just a sense of some of the identifiers that are listed in the protocol parameter registries.

We talk about IP, Internet Protocol. There's things like header flags, port numbers, type-of-service values. When we send emails, the encoding of those emails is encoded in something called a MIME or a media type so you have various different encodings of emails that need to be defined in some fashion so that when an email is transmitted from

one device to another, there's a common understanding about what encoding that email is in and so on.

So there really are hundreds of these registries that are -- I think protocol parameter registries, in particular, are notable for the fact that as an end user, you probably are not aware of most of them.

There is an index of the registries at iana.org/protocols and you can find a comprehensive list of them there. This is just a very small sampling of them.

So where do these protocol parameter registries come from? Well, the preeminent place where Internet standards are developed is the Internet Engineering Task Force, or IETF.

In the IETF, when someone wants to develop a new Internet Protocol, they develop a document called an Internet draft. This is a document that's circulated for discussion among the engineering community.

Once that document is finalized, it's published as what's called a request for comment or an RFC. Despite the name, an RFC is the finalized version of a particular document that's being developed.

Now, many of those RFCs, when they require a registry, a protocol parameter registry, to be devised, will contain specific guidance to the IANA department on how to do that. There will be a section called "IANA Considerations" and in that "IANA Considerations" section, it will tell the IANA what registries are required to support operation of this protocol, what the registration policies should be, i.e., what are the rules IANA should follow in allowing new entries to be added to the

registry, and then what the initial registration and reserve values are for that registry.

So how do we interact with that RFC document process?

It starts before the RFC is even issued. We're involved in a review process as the documents are finalized. And then after the RFC is published, we then have an enduring obligation to implement and maintain those instructions that are inside the RFC.

So let's drill down into that a little bit.

Here is an actual cut-and-paste from an Internet draft. This is the kind of instructions that we receive from the IETF in their standards documents.

So as you can see, it explains we need to create a registry. It tells us how that registry needs to be devised, what it should be called. It tells you that, for example, registrations should be allocated on a first come, first served basis. Registration requests must include a key. I mean, there's various different specific aspects of this registry that IANA is required to follow.

Now, every registry has different requirements and there's thousands of different RFCs. So for each different unique purpose, there's a different set of rules on how the allocations are made.

These are documented inside the RFCs.

So what happens is, when a document is just prior to being published, the RFC editor contacts IANA, the IANA department, asks us to review it. We make sure that we understand the instructions clearly. If there's a

lack of clarity, we go back, we speak to the authors of the document. We make sure there's a common understanding so when the RFC is published, IANA can accurately implement the wishes of the authors and what they're trying to accomplish in the standard.

So an RFC is completed. It's approved. We now have a new Internet Protocol or a new standard.

ICANN's IANA department then creates the new registries required by the document or some standards augment existing registries, so we would then update the existing registries to match with what the document requires.

We then maintain, on an ongoing basis, this registry. We accept registration requests. Some kinds of registries allow anyone off the street to come to IANA directly and ask us for a new allocation to be made in that registry. If that's allowed per the registry's rules, then we -- we maintain that process.

Alternatively, other registries require some formal standards mechanism for future updates.

So following the rules of the registry, the registration policy, we maintain and update the registry moving forward.

So if you go to that Web site I mentioned a few slides ago, you'll find there's actually over 2,800 different registries we maintain.

Again, 2,700 or more of these you've probably never heard of. That's the nature of the business. I mean, a lot of them are very esoteric protocols, but some are critical.

Again, the common element is, generally speaking these protocol registries are behind the scenes. When you interact with a Web browser, with email and so on, you don't see these numbers directly but they are crucial to interoperability of the Internet.

So how do we process these?

Well, firstly, we receive a request. We identify what is actually being asked of. We look at the registration policy. We look at what are the rules governing the type of parameter that's being requested.

I mentioned earlier that the defining RFC determines what that policy is.

We then process and evaluate, following that policy.

A lot of our registries are highly specialized, so we will often go to a designated expert on that particular topic that can help us out in evaluating such a request. I know I can see at least one in the room that is involved in that process with his hand up.

So, you know, IANA doesn't profess to be a technical expert on every single one of these registries. When there's a highly specialized evaluation that needs to happen, we rely on the experts in the community to help us in that assessment.

If there's more information needed, we go back to the requester, but once that processing evaluation is completed, we then make the updates to the registry and then notify the requester that the request has been complete.

So just as an example, how does a request for a new media type work?

So let's assume that you're a software implementer. You've invented a new file format. You know, your new technology has a new way of storing data. It's a file. It's a new extension, let's say. How do you make sure that when that's transmitted over the Internet, it's understood correctly as a file attachment in a Web browser and so on?

What you need to do is you need to come to IANA and apply for a new media type.

So we receive such requests on a form. The form has a number of questions. We review the request. We make sure all the fields have been completed. We send it to a technical expert designated for media type registrations, make sure they find that the request is acceptable.

We add the media type to the registry.

We confirm with the requester.

Process complete.

So generally speaking, that's kind of the workflow in most cases, that we receive an application, we assess it against the criteria, consult with experts if needed, complete the request, update the public registry -- everything we do is public -- so everyone can now look at this list and see that that new file type has been issued, here is the media type that can be implemented in software.

So how many of these do we do? We do around 4,000 of these requests per year. Generally, it's pretty stable. Most of these requests are from the implementer community. They're not from end users,

they're from people implementing software. Software developers and so on.

How do we do this in terms of performance?

Well, in the beginning, I mentioned the genesis of all these registries is essentially the IETF community. So the performance standards on how IANA should operate and how quickly we should operate are defined in conjunction with the IETF.

There's a memorandum of understanding between ICANN and the IETF in terms of what the performance metrics should be for doing these protocol assignments and how that should be monitored.

So beginning in 2007, we started reporting on service level agreement deliverables on a monthly basis, and then as an annual process, ICANN meets with IETF leadership to review the SLA from the previous year, make any amendments that are necessary to govern the next year's operations.

So you can see at the bottom this is the full SLA tally for the last calendar year. We met all of our SLAs for 2014. As you can see, in terms of the KPIs, we met them between 98 and a hundred percent of the time, month in month out.

So with that, I'm going to hand back to Elise, who will tackle the second section, which is number resources.

ELISE GERICH: Thank you, Kim.

I did just want to mention at the end of the protocol parameters piece that those are the largest numbers of requests we receive. They're probably the least understood or known by the general community, but most of our work is done with protocol parameters. Those are the largest number of requests that we receive and process.

So back to numbers.

And just a little graphic to show the three numbers that are highlighted in yellow. The three types, I guess, of categories. IPv4 addresses, IPv6 addresses, and AS numbers. "AS" stands for "autonomous system."

Those are the numbers that we allocate based on global policies from the regional Internet registries.

So all of these allocations that are done of the numbers -- IPv4, IPv6, and autonomous system numbers -- have deterministic formulas that determine -- that's why it's called "deterministic" -- how they are allocated from the IANA department at ICANN to a regional registry.

So the IPv4 addresses are allocated based on a schedule now.

As you've probably heard, it's been in the news for years and years that, you know, we're running out of IPv4 addresses. Well, ICANN ran out of them in 2011. So we allocated the last big block of IPv4 addresses to all the regional Internet registries in 2011.

All we have now are little bits and pieces that have been returned, which we hold onto, and there is a deterministic formula which takes care of how those are allocated.

So we don't really get requests for IPv4 addresses. All those requests go to the regional Internet registries.

IPv6 and AS numbers are allocated on receipt of a request from a regional Internet registry.

We deal only with the regional Internet registries. End users are referred to their local region. So for instance, if someone from Singapore were to ask us for an IPv6 address, we would refer them to the APNIC, which is the regional Internet registry for this region.

So just to reiterate, for the allocation types, for IPv6 and autonomous system numbers, basically there's a formula that allows someone to generate a request from the regional Internet registry to us. So if you are within Singapore, say, and you're building a network and you want an autonomous system number, which is necessary for your network to talk to another network, you would go to APNIC and they would have a formula that would say, "Okay, you can get this autonomous system number."

If the APNIC had run out of autonomous system numbers, they'd had a pool at some point in time and now they're down to just a percentage of that pool, then they could come back to us at the IANA department and say, "I've met the criteria. The pool of autonomous system numbers that I have now is below the formula. Can I have more autonomous system numbers to serve my region?"

So that's how IPv6 and autonomous system numbers work with the formulas and the requests.

As I mentioned in the previous slide, for IPv4, we really don't have a huge pool at the IANA department of IPv4 anymore. Just bits and pieces. And so there's a schedule and a formula that was created by the regional Internet registries, and it became a global policy, and so that's how those allocations are done.

And then there are some special numbers, because most people don't know about them, just as they don't really know much about the protocol parameters, that are done by the IETF, and these are special number allocations for research and multicast and some other types of special allocations.

Those are defined within the IETF protocol specs, the RFCs, and we have a pool that have been allocated and reserved for use by the IETF. So they do not go into the general population.

So just to walk through the steps of how we do this, first of all, we receive -- "we" being the IANA department, will receive a request from a regional Internet registry. That's what an RIR is, regional Internet registry. We have to look and see if they qualify. If they do qualify, this is an IPv6 request. If they have less than a half of their /12 -- and a /12 is an enormous amount of numbers. I think we handed out /12s to each of the RIRs six years ago, and no one has come back for any more IPv6. So this is the criteria on which they can send a request to the IANA.

And then when they send that request, as you can see, there's a formula that determines what they get. So we use the data from the RIRs as to how many addresses they've assigned or allocated out of their /12s and what's available. And this is an example from the RIPE

NCC IPv6 pool. RIPE is the regional Internet registry for the European region.

I think people are starting to peel off and leave this presentation so I guess it is less interesting than I thought. I'm sorry. I was hoping you were out the door before I said that.

[Laughter]

Okay. So this is an example of what the table looks like that we maintain. So basically it says what the prefix is for an IPv6 address. No one knows how to read these and say them in real English, so it is not like IPv4 where you just can say 122.1.1.1. They are more difficult to say.

Then it says who is holding that address block, the date it was assigned, and the status. This is pretty much the type of table that's maintained.

So we communicate the allocation to the regional Internet registry. And once they receive it, they communicate it to their operations community in their own region. And back to the IPv4, what happens there? So based on the formula, there's two allocations of what's called the returned IPv4 pool. That means addresses that were returned and were no longer used by whoever had them originally.

And so we have a formula, which if you have really, really good eyesight you can read. It is on the slide, and it is posted online so you can also download it and you can run a prototype against the pool to see what the next allocation will look like if we don't receive any more returned addresses.

And this formula runs against the returned pool of IPv4 addresses twice a year, once on March 1st and once on September 1st. So it is not a request that comes into the IANA department, and the IANA specialists don't have to process it. Basically, the formula runs. It is on a schedule. And as long as there are addresses in the returned pool, we continue to run the formula against that pool. And then we communicate the results through the registry that we maintain and to the regional Internet registries who received the addresses.

People get all excited about Internet addresses. And, basically, as you can see over the last four years, these are the number of allocations that the IANA department has made to the regional Internet registries.

Obviously, a lot of this has to do with the depletion of IPv4 addresses and the fact that IPv6 is such a huge space that the regional Internet registries have not yet had to generate any new requests for IPv6 addresses.

So as Kim mentioned with the protocol parameters, we do have some performance targets and KPIs. As you can see, we have all these green checks. You may not be able to read it, but it is online at iana.org/performance. And we post these monthly.

However, it's a lot easier to get a green check when you don't have any requests.

[Laughter]

So, basically, everything we've processed, the five requests we've processed in the last year, we processed within the SLA. But I'm still proud to say we processed everything within the SLA, even though

there were only five requests. And my colleagues over here are laughing at me from the regional Internet registries. But we work very closely with them on the establishment of global policies that come through the regional Internet registries and make sure that we implement, for instance, that IPv4 formula for the returned pool as they had intended it to be implemented.

So that's all I have to say about the allocations and assignment of IPv4, IPv6, and the autonomous system numbers, and the registries that we maintain on behalf of the community.

I will now pass it back to Kim Davies, who will talk about domain names and root zone changes.

KIM DAVIES: Thanks, Elise.

I think it's no surprise looking at the agenda of an ICANN meeting that this is perhaps the most interesting part of what IANA does.

So I mentioned earlier that we have all these protocol parameters that are dished out based on standards developed at the IETF.

Before I go further, it's worth noting that the Domain Name System itself relies upon a number of different protocol parameters you might not have heard of. There's things called record types, DNS security algorithm types, DNS header flags. Again, for the DNS to interoperate, these need to be assigned in a cohesive global way. Again, like I mentioned earlier, these are allocated to software implementers directly. It is not something end users come to us to request. But,

nonetheless, that is a fragment of what we do for the Domain Name System managed under the IETF processes and IETF rules that are developed there.

But when we talk about the domain names at IANA, typically what we're referring to is not those components but the namespace, the domain names themselves, the TLDs, the second-level domains, and so on. So that's one aspect of the unique identifier space and use of the Domain Name System and typically what we refer to as domain names.

So this is probably not a surprise to many, but the domain name space is a hierarchical system. You can think of it as a tree. The top of the tree is something called the root of the domain name space. Underneath the root is top-level domains. Under those top-level domains are second-level domains, third-level domains, and so on. This is a unique property of both number resources and domain names. Number resources similarly have sort of a root -- the IANA root. They're divided into chunks. They are divided into smaller chunks. The domain name space works in a similar way.

So -- actually, I will skip back. So the key point to this slide is actually that the root is the primary responsibility of the IANA in the context of the Domain Name System. Maintaining the root implies maintaining delegations of top-level domains. We maintain the database, the official record of which top-level domains exist, who they are operated by, and so on. And I will drill down into how that works moving forward.

So why would someone come to IANA to make a change to the root? Well, generally speaking, an event happens that triggers a request. So,

firstly, there is a change in the TLD operator. This can be the creation of a new TLD, so you are going from no operator at all to a new operator. Or it is taking an existing top-level domain that needs to be transferred to another entity.

Another kind of change that triggers a change to the root is just routine maintenance. Maybe it is a technical update. The technical infrastructure for a particular TLD needs to be changed to support new servers or a new environment. Or maybe it is a staffing change. Maybe the people that work at the TLD operator have changed and our contact point needs to be updated with that TLD operator.

Another scenario, less regularly thankfully, maybe there is a natural disaster. We have these Country Code Top Level Domains based in every country of the world. Countries sometimes have a hurricane that knocks significant local infrastructure offline, earthquakes, that kind of thing. Often that will necessitate some kind of emergency change request for the TLD.

So this is the actual information that IANA maintains on a top-level domain, on a TLD. We maintain the record of who the recognized company or organization is that operates the TLD, its formal legal name, its physical address. We maintain two contacts, an administrative contact and a technical contact. These contact persons, we maintain their name, job title, company address, phone, fax, email.

We maintain the technical configuration. This is the data that enables the TLD to actually function. It is a list of the name servers that run that TLD. It is the I.P. addresses of those name servers, and it is the DNSSEC records that support security for that domain.

And then we also list some additional information, some metadata. It's not critical to the operation of the root zone, but it's nonetheless very useful for the community, things like what's the Web address for that particular registry? Where is their WHOIS service located and so on and so on.

So taking that from the abstract to the actual -- this is an actual cut and paste of the data that we maintain for .HAMBURG just as an illustrative example. Again, contact persons for that domain, this is all public information. None of this is secret. It is all posted on the IANA Web site, available in the IANA WHOIS service and so on.

Okay. So we've received -- we've had a triggering event that necessitates one of those changes.

The TLD operator submits a change request to us to process. We have an automated online system called RZMS. It is the root zone management system. Typically TLD operators will use this system to lodge a request, even though they can call us, fax us, and so on. They submit a request. It is a form. They spell out in detail the changes that they want to have made and submit it through this process.

So we receive a request. Now much like the IETF protocol parameters and much like the number resource allocations, there's policies, there's procedures. So we review the requests. Make sure it agrees with the policies and procedures that are in place. Some of them are technical. Some of them involve consent. If there are any issues found, we clarify with the TLD operator. And then once we've clarified and we believe the process should proceed, we forward the request to NTIA which I will get to in a moment.

Just drilling down a little bit into what the policy checks we do are, I mentioned some are technical. There is a set of name servers that are supposed to operate for the TLD. We check them all and make sure they actually work. We make sure that when we send requests to those name servers, they return the correct data that matches the request that they've submitted.

We verify that the DNS data that we receive matches, the DNSSEC records match, and so on and so on and so on.

If there is email addresses in the submission, if it is a new email address, we send an automated email to that email address with a link they need to click. That allows us to verify that the email address is functional as well.

The next area we look at is consent. We need to make sure that a change request is approved by the appropriate people. So anyone can submit a change request. But on behalf of the existing TLD operator, they need to consent to any updates to their record. So we reach out to the existing contacts we have on file and make sure they agree.

If the request involves implementing a new contact, we reach out to the new contact and ensure they agree to being a contact.

Certain types of specific technical changes can impact multiple TLDs. In those specific cases, we reach out to the other TLDs that might be impacted by that change and make sure that they agree as well.

We make sure the request meets any legal requirements, regulatory requirements. We make sure that the supplied data is clear, well-formed and consistent. And our job is to maintain a database that is

easy to be used by the community. So if the material they've supplied for us to list publicly is not clear, we go back and seek clarification.

And then specifically in the case of transferring or, as we say, redelegating a TLD from one operator to another, there is an additional set of policy requirements that I will get to.

I will get to right now. In terms of the policy of doing a transfer, it is really easily divided between gTLDs and ccTLDs. GTLDs are governed by GNSO policies. So what we do is if a transfer is to be conducted between one gTLD and another, we essentially look for a contracting process to have been completed. If such a change was to occur, the new operator would have had to sign a contract with ICANN to effect that change.

So we look for the completed contracting process with the other departments in ICANN. Once that process is complete, a transfer can be conducted within IANA.

And this is very different from what we do with ccTLDs, these Country Code Top Level Domains. What we are looking for there, in essence, is that some kind of consensus-building process has happened within the country. There is a process of asking for materials to be supplied concerning that consensus-building process. But the essential element is that it is not an ICANN decision about who should operate a ccTLD. We're looking for evidence that that decision has been made inside the country before we conduct a transfer process.

So we've gone through the validation process. We are now submitted for verification by NTIA. I know it is the topic of the year in the ICANN

community. This may not always be the case moving forward. But for now, as we receive root zone change requests, once we as ICANN have done the validation, we submit to NTIA for verification. They look at the change. And once they're satisfied the change has been implemented correctly, been reviewed correctly, I should say, they give authorization to proceed.

Once we receive that authorization, their process is implemented -- sorry, the change is implemented. Technical changes are implemented in the root zone. This is done by VeriSign.

A tamper evident seal is applied to the root zone using DNSSEC. A distributed version of the new root zone is distributed to the root server operators. And then that root zone database that lists the other non-root zone material, contact phone numbers and that kind of thing, that's updated based on the authorization.

So that's kind of the process. But that's not all we do in terms of managing the root zone. Another key piece of managing the root zone is maintaining the security component, the DNSSEC component of the root zone. So we have this thing called the root key signing key. And as part of our obligations for managing root zone-related functions, we maintain this key signing key. It is kind of the apex of the trust system that's used to secure the DNS.

And the way we do this is we have this auditable process of performing key signing ceremonies. We have this key stored in a very secure way, and we conduct ceremonies every few months where we retrieve this key, perform technical operations, and return it back to its safe environment.

And the way we do this is we involve members of the community as key participants in this process and I will walk through exactly how that works now.

So we have this key. It is actually like a file, but this key is very important to be kept secure. So the way we store this file in a secure way is in a specialized device known as a hardware security module. Its sole purpose is to store cryptographic keys in a secure way. It is designed to be tamper-proof. So if there is an attempt to open it, the contents will self-destruct. So if you grab one of these devices, which is, you know, a bit smaller than a laptop, shake it very hard vigorously, it will self-destruct. Drop it, it will self-destruct. Try to open it, it will self-destruct. Set fire to it, it will self-destruct. It has a lot of protections built into it, so any attempts to tamper with the device, it will result in itself self-destructing.

Now, the way this box works is you have a set of smart cards and you need to insert a set of smart cards to switch the thing on.

The way the device we use is configured is there are seven smart cards for the device, and in order to use the device, three of those seven smart cards must be present to switch the device on.

So how do we protect it with those seven smart cards?

Well, we have this system set up where we give one of each of those seven smart cards to a member of the ICANN community. These are members of the ICANN community from the sort of technical side. They volunteered when we set up this system. We call them trusted community representatives. The idea is we give one of those cards to

each one of these community members. In order for us to use this device, we need to get three of these community members together, conduct a ceremony, in order to switch on the device.

So we have this device that's protected with smart cards. People that need to be present. We take this device and we put it in a safe. This safe, the combination is only known by special individuals known as safe security controllers. This safe is protected by seismic controls, so if there's an attempt to open the safe, we're aware of it.

There's other sensors designed to identify surreptitious entry into the safe.

And it's a high-rated safe as well, so if you wanted to try and drill into it, it will theoretically take hours and hours and hours of persistence in order to be able to pull that off.

So we have the HSM. It's in a safe. We then store it in a safe room, a secure room, that can only open -- be opened jointly by two different people. We call this the ceremony administrator and the internal witness.

This room itself is monitored with intrusion detection systems, motion sensors, and so on.

So you can see I'm painting a picture of multiple levels of control of access to this device.

So far, you need to have these two people that can get into the room. You need a different person needed to open the safe. Then you get the

device and you need three of those seven ICANN community members to be able to switch the device on.

So we take that room, we put it in another room. This is our ceremony room. This is where we conduct our ceremonies where we access this device and generate the keys needed to secure the DNS.

Now, access to this room is controlled by another person called the physical access control manager who is not on-site. It's someone that's remote. They need to preauthorize access to the room in order to be able to get into it.

All accesses to the room are recorded on video. When we're in there, it's recorded by participants. We're audited by PricewaterhouseCoopers, so they do a third-party audit. They're in the room as we do the ceremonies, to make sure we're following correct procedure.

And so wrap that all up and then we have two duplicates of that system. We have one on the U.S. West Coast in Los Angeles -- specifically El Segundo, California -- and we have one on the East Coast in rural Virginia.

These are designed to be exact duplicates so that if one is unavailable, we can use the other, but in practice, we tend to alternate between the two in our regular operations.

So we have these facilities. We need to get access to this HSM to secure the DNS on a regular basis, so four times a year we conduct a ceremony.

A ceremony is where we take three of those seven trusted community representatives and others to one of these facilities and we use these HSMs to sign DNSSEC keys that will be used for the root zone.

This process is streamed on line. It's recorded. You can look at all the videos in they're breathtaking glory on our Web site. There are hours and hours and hours of videotape there. It's all on line at iana.org/dnssec.

But I guess the key point here is all the protections I just mentioned and involving community members and doing it in this open and transparent way is designed to ensure trust in the process. The idea of DNSSEC hinges on the notion that that key inside the HSM is not compromised, so we have these multiple layers of security to make sure that the community trusts that that key is only used for its intended purpose.

So here's a picture of all the participants in the last ceremony we did, which is only about two weeks ago. If you're interested in this topic, we actually had documentary crews film it last year twice. We had The Guardian newspaper come. We also had BBC's Horizon show come as well. So there are some links there, and you can watch some videos of it in action, to see how it works.

That's a summary of how we manage the keys for the root zone, but I think it's worth noting that, you know, security for us at IANA is not just DNSSEC, it's one of the principles by which we do all the functions we've talked about today.

We have dedicated workflow systems that are unique to the IANA functions. We don't use a system that's used across the whole

company. We have dedicated IANA-specific systems. And access to those systems is limited to IANA staff.

So the IANA systems are managed by the staff that Elise identified at the beginning of the presentation.

We have a separation of user-facing and staff-facing systems, so the system -- the root zone management system I mentioned earlier, for example, there's a system that customers log into. That is separate from the system that staff logs into to perform activity. And the system that staff logs into has a higher set of controls in relation to getting access to that system.

Importantly, we have regular third-party audits. I mentioned earlier that PwC does an audit of those DNSSEC ceremonies. They also do another kind of audit called a SOC2 audit, and what they do is they look at all of the controls we implement in terms of security of our internal systems for IANA. Things like our password policies, our security controls and so on. And they audit to make sure that we're following all of those procedures. And in fact, we're hoping to wrap up our second year of doing it later today.

So we've -- we did it for the year 2013. We're now wrapping up an audit for the year 2014.

In terms of performance, jumping back to root zone management in general, the number of requests we receive varies from month to month. 123, 77, 81, so on.

The KPIs, you know, are not as impressive as the ones we saw earlier, and the reason for this is that root zone management is -- is a more complex beast, let's say.

There's often a lot of back-and-forth between the customer and us. There's policy processes that need to be engaged. And the way we've designed our KPIs is we monitor end-to-end performance.

So this includes time that is spent with the customer clarifying information. It's time spent with the root zone management partners implementing change requests.

So, you know, it's a little -- it's a more aggressive KPI, so we're not, you know, a hundred percent, 99% all of the time.

But we felt, and in consultation with the community when we developed these, it's a more fair assessment to measure end-to-end processing time rather than just limiting ourselves to the time it takes ICANN staff to do change requests.

So there was one month in the reporting period where we didn't meet all of our SLAs. You know, this was simply triggered by a very complex request that went for eight months, but that request alone made us fail one of the SLAs that we report on.

All the SLAs, all the details of the SLAs, is posted on our Web site, iana.org/performance. You can go there and look down and drill down into the detail of the change requests.

So with that, I'll pass it back to Elise for the closing segment.

>>ELISE GERICH: Thank you, Kim. And I do notice that more and more people peeled out, so I guess it is truly boring to listen to exactly what we do on a day-to-day basis, but it's an important job and the staff, the 12 folks that I showed you at the beginning, take it very seriously and work very hard to meet the needs of the community overall.

So how big is the job that we work on? And this bubble chart tries to represent in total the types of things we've done from January to December in the calendar year of 2014.

And you can read the charts. I don't think I need to say the numbers out loud. But as you can see --

We'll have a Q&A at -- in a minute. Okay. Thank you.

We have quite a few requests, and most of them are in the protocol parameters area, the area that most people know the least about, and which are really the foundation for all of the Internet activities that we all know and love today.

We have a lot of domain name requests and things that go into the root.

And the numbers, which is quite a critical function, the RIRs and ICANN have a very good system but there are very few requests. But we work quite closely together with that community also.

And we have third-party audits and some general inquiries.

It's amazing. Some people think we're the, like, customer service for fixing your laptop. It's -- we get some very interesting general inquiries that have nothing to do with the IANA functions, and some that do.

So this is the fourth year we've done a customer satisfaction survey, and we broke it down by customer group, so as Kim explained earlier, that trusted community representatives is the small pool of people who help us with the DNSSEC key signing ceremonies. The regional Internet registries, you know. We also do changes to the .INT. That's the top-level domain for international treaty parties. And the IESG. And for those of you who don't speak IETF-speak, that's the steering group within the IETF, the Internet Engineering Task Force.

And as you can see, we got very high marks. We were quite pleased that the communities felt that we were generally meeting their requirements and needs, and we are working on the little red areas to see if we can convert those people to "very satisfied" or -- or "satisfied" in next year's survey.

And since domain names is really the topic that seems to get everyone engaged, we have put up here an example of a dashboard that we publish monthly on the iana.org Web site that kind of shows you the numbers of requests we get in the domain name root zone space and the average processing time.

Finally, a lot of people think we do things that we don't do, so we've tried to kind of clarify in a very high-level bullet perspective what the IANA department at ICANN does and what it doesn't do.

And so I'll just focus on what we don't do, because that seems to be where all the confusion is.

We don't create nor interpret policy. The policies that we implement come out of the community or out of -- or the RFCs. And the RFCs are

membership -- not membership-based. Community-based. And those are done by the community and then we implement them.

And we don't determine what is a domain name. So if you wanted a country code top-level domain, Jon Postel in his wisdom said all country codes would be defined by ISO-3166, so if you happen to be in the ISO-3166 table as a country, you can apply for a country code top-level domain.

And we don't determine who can be the manager of a TLD, whether it's a gTLD, a generic top-level domain, or a ccTLD, a country code top-level domain. That's done for country codes locally within the community, so the country, and for gTLDs that's done through the gTLD process that came out of the GNSO.

And in summary, what we do as the IANA functions operator is to maintain registries of unique numbering systems. Those are the unique systems for the protocol parameters, the domain name system, and the Internet addressing systems.

And most of these registries are very straightforward. It's an administrative task to make those happen. We do have to implement new registries based on the criteria that are defined by the policymaking group that come up with the registries.

And these are used by the world in general.

And most people don't even know how many of these registries are used in their day-to-day Web surfing or their refrigerator, even, if your refrigerator is communicating back to some service group. A lot of

online appliances today have all these protocols that we've talked about embedded in them.

So I just wanted to conclude with the fact that we're really glad that so many of you were interested in finding out what the day-to-day job is of the IANA functions operator, and we'll open this for questions to Kim and me, and I think, Carole, you had a remote participation question, to start.

REMOTE INTERVENTION: Thank you, Elise.

The question comes from Dan Rogers. Is there any plans to reclaim IPVs which are currently issued in large blocks but are unused?

ELISE GERICH: I'll respond to that.

There have been activities over the years to try and reclaim what are perceived to be unused IPv4 address blocks. It's not always obvious what's used or unused because a lot of large blocks are used for internal networks. And I know the RIRs in their own regions have looked, over time, at trying to reclaim address space because they're closer to the customers themselves. The users.

And those have been fairly futile activities over the last decade to do that, and at this point in time, I think all the energies are focused to transitioning to IPv6.

There's so few IPv4 addresses, in general, that what you really should be focused on is where the future's going and not try to continue to come up with conservation, I guess, of IPv4. It's pretty much done. You know, the end of scarcity is here, if you use IPv6, and if you continue down the path of just trying to reclaim IPv4, it's -- it's kind of like, you know, having one grain of rice and trying to feed a thousand people.

I guess I should say please come to the microphones.

[Laughter]

SIVASUBRAMANIAN MUTHUSAMY: Okay. Sivasubramanian Muthusamy from ISOC India, Chennai, and there was a slide five or six slides before that was on performance, IANA performance. Is it possible to go back five or six slides?

ELISE GERICH: In the domain name section? This one?

SIVASUBRAMANIAN MUTHUSAMY: On the first one, yes. So --

ELISE GERICH: Okay.

KIM DAVIES: Is this it or --

SIVASUBRAMANIAN MUTHUSAMY: There was one that talked about failure rates, some 84% success rate and...

ELISE GERICH: Is it in the protocol parameters or in the domain --

SIVASUBRAMANIAN MUTHUSAMY: Towards the end of your presentation, there was a slide that said the success rate of IANA performance, key performance indicators, and all that.

KIM DAVIES: I think it's this one, correct?

SIVASUBRAMANIAN MUTHUSAMY: Oh, yes. Okay. Okay. And here, it says key performance indicators met 84%. What does it mean? Does it mean that 16% of the times it was not met, and what are the implications? What does it mean when KPI is not met? Isn't it supposed to be a hundred percent?

KIM DAVIES: So the -- the KPI that we've used for the slide -- and again, the URL is there, if you wish to go into the specifics, but we picked the timeliness metric, and I think off the top of my head the way the timeliness metric is written is that we will conduct I think 80% of change requests within 14 days or something like that.

SIVASUBRAMANIAN MUTHUSAMY: Okay.

KIM DAVIES: So in this case, you know, in January, in 84 percent of the time we met that timeliness metric.

You know, we devised these KPIs in consultation with the domain name community back in 2013, and the thinking was that --

The way we process domain name requests is actually quite interesting.

The majority of the requests we do are routine. We receive a change request, it's a routine maintenance update, it's done within a couple of days.

But then on the other end of the spectrum, we have some requests that last for years.

In the case of some country code top-level domains, perhaps there's an internal fight over who should operate the ccTLD and there's a long arduous multiyear process to do that change request.

So we have this ticket open that goes on for a long period. It's not that we're processing it; it's that the community is doing its resolution process.

So this metric captures all of that.

SIVASUBRAMANIAN MUTHUSAMY: Okay.

KIM DAVIES: So the reason we don't, like, judge ourselves based on a hundred percent but rather sort of a threshold of about 80%, is to recognize the long tail of a very small number, but nonetheless requests that last through long periods of time.

SIVASUBRAMANIAN MUTHUSAMY: So it's entirely about process delays, nothing about misdirection or directing to another zone or some technical failure?

KIM DAVIES: I mean, I don't think it's for technical failures or things inside our control, but that's a dialogue we have with our community. You know, particularly with the TLD operators, we report to them more directly, and we have a dialogue about "Here's where we failed, here's what we saw." Yesterday, at the ccNSO, we explained in more detail why there's a red "X" right there, so that TLD managers are comfortable with how that works.

SIVASUBRAMANIAN MUTHUSAMY: Thank you.

UNKNOWN SPEAKER: Thank you for the informative presentation. Two years back, I remember that you received a request to take down the .SY TLD for Syria. I just want to ask who sent that request, if you're able to comment on that. And in countries of instability, how do you -- do you -- is there a process for redelegation or taking down a TLD?

I commend you for not taking down the .SY TLD, but I want to know the process of having those requests dealt with, and if they request -- the .SY request of redelegation came from a Syrian entity, another Syrian entity, or a foreign government.

KIM DAVIES: I can respond. So I honestly don't know if there was such a request. I don't recall such a request. But conceptually if there was such a request, the very first step is to ask the designated administrative and technical contacts for that domain to agree. And if they don't agree, we simply don't proceed. And I think .SY is operated by a government agency in Syria. So if we receive such a request, first step is to ask them do you agree to this request. And if they say no, that's the end of it. Does that answer your question?

UNKNOWN SPEAKER: Yeah, thank you.

WOLF-ULRICH KNOBEN: Thank you. My name is Wolf-Ulrich Knoben. I'm with the ICG. Thank you very much for this comprehensive information. It was not boring. It was very interesting, especially the last part of the keys, which is a very interesting thing. Thank you.

I have a question. Can you share with us some information about the -- kind of the -- well, the breakdown of staff which is allocated to all these tasks and in terms of these three task areas. That's the first question. And maybe as well of budget allocated to those tasks.

And the second question is, as we are talking in light of the discussion of the transition, the question is: Is it from a management point of view possible to break up those tasks, to separate it from each other. Because we have these discussions over there the other day as well that come up in the future with maybe three separate IANAs, allocated protocols, numbers, and names. And what would be the implications to do so?

ELISE GERICH: So, Wolf-Ulrich, I will try to take that one. And if I forget a question, I will come back to you. One of the first ones is what's the breakdown of staff that process requests. As I mentioned, there is 12 of us total and four people are dedicated to processing the requests. So that's basically four divided into 12, that's 1/3 of our staff handle as the first tier processing of the requests.

Kim and Michelle, Davies, and I are the specialists for the protocol parameters area, the root zone and domain name area, and the regional Internet registry area. So we're kind of the second tier of customer support there.

But for the actual processing of the requests, four of our staff members are dedicated to that.

Then you've asked a question about the budget. ICANN publishes the budgets, and the IANA budget is published. So you know in total how much is in our budget, and about 50% of that budget is personnel costs.

And then, finally, I think you asked: Would it be possible organizationally to separate so that each function had its own set of

people? When I started at ICANN about five years ago, we were much more siloed like that. We had two people that focused on one area, function. We had two people that focused on another function and one person who focused on the third function.

It became quite evident that those -- that number of people are perfectly capable of handling all the requests. However, if people go on vacation or get sick, there's really no backup and there's no redundancy. So part of my activities when I first joined ICANN was to create a team that could back up each other so that we have a pool of people that manage all the requests or can manage. They are specialized. For instance, some people specialize in the protocol parameters area, but they're also trained for the numbers and the root zone area and vice versa. We have people that are specialized in the root zone area, but they are trained to be able to support in the others.

So I think given -- you've seen the number of requests we handle. Having a small team to do those individual functions means that you probably would either be overstaffed and underworked or you'd be understaffed and overworked.

WOLF-ULRICH KNOBEN: Thank you.

ELISE GERICH: You had your hand up. Do you want to come to the mic at the front or are you not going to ask your question? Okay, great, please.

SEUN OJEDEJI: Thank you. Thank you for the presentation. My name is Seun Ojedeji, from Nigeria. And I'm speaking on my behalf. I just want to ask for information if you could explain how these figures come about. That is, do you normally discuss these figures with the respective communities to agree on them before they are then published? Just want to know does IANA staff come about with these figures or it has been agreed among the communities? Thank you.

ELISE GERICH: So let me repeat your question, Seun. I think you asked us if we discuss these figures with the communities, right? And so, yes, part of our responsibility is communicating and reporting to the various communities. So, for instance, with the numbers and the KPIs for the regional Internet registries, we attend and give a presentation and demonstrate those numbers at those meetings for the IETF protocol parameters.

We have meetings at every IETF with the leadership to go over those numbers. And also as I mentioned, Michelle Cotton, who is our liaison and engagement person with the IETF, she has monthly phone calls where she goes over the numbers that are related to the protocol parameters. And we meet at each ICANN meeting with the ccNSO and talk about the performance numbers there.

And then internally, Marilia, who I showed in our picture earlier, is tracking on our performance numbers to see if there are areas for optimization. For example, KPIs for timeliness. If we set the target based on community consultation at 80% of the requests are answered in a certain amount of time, is that a good baseline after all? Should we

be raising the bar on ourselves? If we're always getting 90%, why are we measuring ourselves so lowly? Or if we're never measuring it, what are we doing and how can we fix it? So, yes, we discuss those things with the community.

SEUN OJEDEJI: Thank you. Just to follow up, has there been any instance where you presented a particular figure and the community denied -- I mean, disagreed with you that you actually did not meet that particular figure which you presented?

ELISE GERICH: I don't recall any time when I've been present or any of my staff have been present where people have thought that we've measured the numbers improperly. But thank you.

KIM DAVIES: Just to add to that really briefly, the raw data that's aggregated into these states are actually published as well. So you can drill down into excruciating detail, if you so desire.

ROBERT GUERRA: Hi, good morning. Thank you for a great presentation. My name Robert Guerra. I'm with the SSAC and also on the CWG that's working on the IANA transition.

If you could go back to the slide that talks about the NTIA role. And I think something that I think maybe has not been mentioned and just

what my question will be based on is: A lot of the discussions have been about NTIA relinquishing its role. So when you talked about, I think, there was a step that it goes to NTIA verification, just maybe for -- maybe for the folks here, if you could maybe go into that a little more specifically, how long does that step usually take with NTIA? And what's the process to deal with issues if they arise? Because there is a lot of misunderstanding in the community as to how evolved or not that process step is. Thank you.

ELISE GERICH: So the validation or verification or authorization by NTIA is usually well within 24 hours. I can't tell you if it's within two hours or three hours, but it is well within 24 hours. And, occasionally, it might be as long as two days, but it is rarely any longer than that.

And I'll defer to Ashley Heineman who is with NTIA.

ASHLEY HEINEMAN: Typically it will take us within two to six hours to respond to the request. But when it comes to actually doing the verification and authorization, that actual step is actually a matter of minutes most of the time. So it's not very time consuming.

ROBERT GUERRA: Just as a follow-up. You told us about the number of staff you have at IANA. I'm also just curious. What's the number of staff at NTIA that deal with IANA? Thank you.

ELISE GERICH: I will defer again to Ashley Heineman.

ASHLEY HEINEMAN: We have -- currently associated with the contract overall, we have the primary and that's one person and an alternate, which is myself. And in both cases, this is not a full-time job. It's a couple of hours a week that's associated with it.

ELISE GERICH: Please.

UNKNOWN SPEAKER: Good morning. My name is (saying name), and I'm from Lebanon. And I'm speaking on behalf of myself. I just had a request about the request. If you could give us examples of the types of requests you get for parameters and numbers and names.

ELISE GERICH: So, for instance, an example of a request for a protocol parameter is there's something called the private enterprise numbers pen. And we get hundreds of requests there, and that's individuals from the community that want to have a number assigned to them privately. It's called a private enterprise number.

And so we'll receive that request. We, again, then check to see if they meet the criteria and then we'll assign the number.

Is that the type of example? Okay.

And then from the RIRs, the regional Internet registries, primarily what the types of requests we get now are for autonomous system numbers. Autonomous system numbers are the numbers that are used by ISPs and network service providers so that they can talk to each other. And those are primarily -- we'll get the request. We'll check to make sure that the pool of autonomous system numbers that the regional Internet registry was holding has hit the criteria, it's below a certain level, and then they are qualified to receive another block of autonomous system numbers to serve their communities.

UNKNOWN SPEAKER: For the names?

ELISE GERICH: For the names? Oh, we had that slide if we can pull it back up so I can just -- it defines the different -- we have what are called DS record requests. That's for the DNSSEC change in the record in the root.

Where's the one at the end? No, examples, the chart where it says how many we get in a month. It's at the very end.

So we'll get that. We'll get a request to change the administrative or contact information.

We may get a request to change metadata. Someone's changed their URL for their Web site and they want us to make that kind of change.

Then there's the changes of new gTLD delegation, which means that we will follow the steps to confirm that they've met all the criteria and then they can be added into the root zone.

So the slide behind me has name server requests, that means the servers that maintain the DNS information, DNSSEC records, admin and contact changes, metadata which I just explained, delegation or redelegation requests, and then a root server update. Thank you.

KIM DAVIES: Sorry. Just a time note, we're over time now. So we will draw a line behind those that are already standing up at the mic.

PAUL KANE: Thank you very much. My name is Paul Kane. I'm a CC manager and a member of the CWG working group. Can I thank you both very much for taking the time to come and give this presentation. It is very useful to have the dialogue.

One of the things that has been made very clear today is the process is relatively lightweight and simple. And as part of the transition, I believe we need to keep it simple.

What may help in that task is if you could offline later give us as a community the KPIs that you are currently working to because I think it would be very sensible to try and formulate some form of service level expectation so that you know what you're currently doing, we as a community can review that, see if it meets our needs, which incidentally it does. As our registry manager, I'm very happy with our performance. But the emphasis needs to be on keeping the process simple.

We learned NTIA just have two hours a week. You have 1 1/2 staff-ish doing the job. If we were to automate it, it makes it more efficient. We

need to keep it lightweight and efficient. If you wouldn't mind providing the parameters under which you develop the KPIs so we can develop a service level framework, that would be helpful. Thank you.

ELISE GERICH: We're happy to do that with whomever would like it. And besides -- if you do go online to the iana.org/performance, before each of those charts where we say we've met service level agreements, we define what the targets are and we also achieved those -- not achieved, we established those targets based on the customer consultation process and the feedback from the communities. So I don't know if we should talk with you or to someone else, but we're happy to --

PAUL KANE: One document, very simple.

ELISE GERICH: Okay, thank you.

Please.

AHMED EISA: Ahmed Eisa from Sudan. My question is to NTIA. Do you think by the end of this 52 -- ICANN 52 meeting that we'll have a suggested names for the transition that will take place, too, for example, like ICANN?

Also, my other questions: Do you think NTIA will accept the transition to take place to a multistakeholder outside USA?

ELISE GERICH: I will defer to Ashley. I think you've asked this question in two other forums that I've been in. Haven't you asked this question before? Yes, okay. We'll see. Now it's Ashley's turn.

ASHLEY HEINEMAN: I might need clarification to answer the first one. But to answer the second one, assuming I understand it correctly, is we are fully committed to move forward with this transition. And we are relying upon the community to present us with a proposal. And assuming that proposal meets the initial criteria that we provided and has a broad community support, we are prepared to proceed. But I will need clarification on that first question, please.

ELISE GERICH: I think Kim closed the line. Kim, can we be flexible?

KIM DAVIES: As long as another session doesn't walk in, I guess we can.

ELISE GERICH: Please go.

JAN SCHOLTE: You haven't heard how complicated my question is yet. my name is Jan Scholte. I'm at the University of Gothenburg. We are acting as an accountability advisor in the transition process.

My question -- Thanks very much actually. This has been hugely clarifying. My question was a little bit on the NTIA thing, again. Three things.

What -- when you verify something with NTIA, what actually is being verified? That's the first question.

And second thing is: What would be a problem in that verification? What would be an example of a problem that takes the two or more days? And the third is, after a transition, who will be doing that verification? Is it still NTIA or some new entity? Thank you.

ELISE GERICH: So I will take a stab at this, and then I will let Ashley come in and tell me if I got it wrong or right.

But from the perspective of ICANN, the verification includes, let's say, the most complex root zone change that they deal with is a redelegation, changing the management of the TLD to a new manager.

There is a number of criteria that are set out in RFC-1591, and the ccNSO has just come up with a framework of interpretation to have more modern guidance on the policies stated there.

And those criteria are the things that we use to develop a recommendation to, yes, transfer the management of a TLD to a new manager.

And so the verification is we write a big report. Kim's very involved in this because he's our specialist for TLDs, and it is a multipage report that then gets sent to the NTIA so that they can see that we have collected

all the data to meet all the criteria and that we have followed the process checking that we have complete information. And so that's the verification.

And once they're content with that, Vernita or Ashley, the two individuals that take care of this, can authorize that the change can take place. Does that answer your question?

And the second question, now I've forgotten what it was.

JAN SCHOLTE: What would be a problem in this process that will actually hold things up?

ELISE GERICH: Well, they've never sent one back. So I think maybe, Ashley, I don't know if after you get your photo taken --

[Laughter]

I don't mean to embarrass you, but the question really is to you. What types of things might cause NTIA to send something back to us?

ASHLEY HEINEMAN: It has become such a streamlined process over the years that, quite honestly, we rely on ICANN essentially certifying that they have done what they are supposed to do.

In the case of a new gTLD delegation, it is basically a template that articulates they have done this, this, this, and that. So it is a very objective process.

Same thing with the ccTLD reporting for the most part, it is very objective. By the time it comes to us, we just make sure that there's nothing abnormal. I can't even point to a case recently where we've had to go back. I don't think that's actually even happened in my ten years, so I couldn't even tell you of an example specifically.

ELISE GERICH: Kim would just like to add one thing.

KIM DAVIES: I think it is worth noting, NTIA actually issued a document a few weeks back clarifying these exact questions. So I would recommend reading that document as well.

JAN SCHOLTE: The last thing, of the transition, who will be doing the verification? Is it still NTIA, or is it somebody else?

KIM DAVIES: I think that that's precisely the question that's being posed to the community and being discussed right now.

Last question.

SIVASUBRAMANIAN MUTHUSAMY: Sivasubramanian. IANA -- and this is a question directed at IANA. You handled this function for so long. What verifications are the situations where the NTIA with its oversight powers stepped in and said, you're not doing something right and you should be doing it like this? Or was there any situation where NTIA came and made non-routine staff changes because it was not performed right? Was there ever a situation where NTIA stepped in and said that you're not handling funds right or your resources right?

UNKNOWN SPEAKER: I'm terribly sorry. We will have to cut this off. The next session is ready to start.

ELISE GERICH: The answer is no.

Thank you very much for your attendance. And I'm sorry, but I guess we're out of time and they want the room. Thank you very much.

[Applause]

[END OF TRANSCRIPTION]