
MONTREAL – IANA – Three Years Since the Transition
Monday, November 4, 2019 – 17:00 to 18:30 EDT
ICANN66 | Montréal, Canada

KIM DAVIES:

Thank you, everyone. Today we have a session that has been put together by the IANA team that we've entitled "IANA, three years since the transition." Recently, we passed that milestone in terms of time, and we felt that it would be a good opportunity to review and recap what has happened since the transition, and how the IANA functions are being operated today. For those of you that don't know me, my name is Kim Davies. I'm the vice-president of IANA services for ICANN, and also president of PTI. To my right is Naela Sarras. She is our operations director for IANA. To my left is Marilia Hirano, who leads our IANA strategic programs for the team.

What happened three years ago? What are we celebrating the anniversary of? For context, from the year 2003-2016 the IANA functions were performed directly by ICANN under a contract with the US Government. Prior to that, in the ancient 1900s, the IANA functions were activities of the US Government under a variety of different other programs.

However, in 2016 something changed. The IANA stewardship transition was completed. This ended the US Government's contractual oversight role and handed responsibility for oversight of the IANA functions to the multi-stakeholder community that's represented here at ICANN.

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

Part of the former model was the creation of a new non-profit organization. We call it today PTI. Formally it stands for Public Technical Identifiers. Others will know it as the Post Transition IANA. Nonetheless, it is a new built-for-purpose non-profit organization that operates the IANA functions today. To explain the diagram there, we celebrated the stroke of midnight on 1st October 2016, when the IANA contract with the US Government ended, by having the whole team out partying. It is a good time to recap, a good opportunity for us to reflect on what's happened, to see what has changed with the IANA service delivery, and also to reflect on what stayed the same. What are the essential elements that continue to operate as expected?

What we're going to do today is do a brief walk-through of what the IANA functions actually are, how they work, and review what the ultimate impact of the transition has been, starting with basic concepts. What are the IANA functions? The best way I have to summarize them is we are the official record keeper for the unique names and numbers that are used by the Internet technologies of the Internet. It is through this record-keeping function that using unique identifiers is consistent around the world. This helps the Internet to work as you expect it to.

The IANA functions predate ICANN. Certainly, they're not an invention of ICANN. In fact, arguably, they're one of the oldest Internet institutions there are. However, there was a recognition in the mid-1990s that the relatively informal oversight model that had existed to that point needed to be replaced by a formal multi-stakeholder system. This led to the creation of ICANN to be the home of the IANA functions.

The unique identifiers, and we'll get into them in a little detail further along, include protocol parameters, Internet number resources, and domain names. The IANA team that maintains these records do so in accordance with policies that have been set by the relevant communities. Internet naming policy is set here at ICANN meetings, predominantly. Number policy is set by the relevant regional Internet registries. The protocol standards communities cover the rest, predominantly in the IETF.

Here on the diagram you can see Jon Postel. He was essentially the founder of the IANA functions, and operated them for many years. He worked closely with other founders of the Internet, including Steve Crocker and Vint Cerf.

Why do the functions exist? Think about what would happen if we didn't have them. If there was no global coordination of the Internet unique identifier systems, then if you typed in a domain name or another kind of identifier, it may end up in an entirely different place, or would not result in successful communication with the other party. Whilst the Internet is free from central coordination for the most part, ensuring all devices are essentially speaking the same language is key to making the Internet work. Those kinds of coordination functions of these unique identifiers is key to what IANA is here to accomplish.

Who uses these registries? All manner of parties. Vendors, service providers, application developers. I think it's fair to say, for the most part, if you're not a software developer the vast majority of what IANA coordinates you will be blissfully unaware of. Predominantly, the

registries we maintain are used by software implementers. They are not seen by end-users of software, but they are critical to the success of software inter-operating. Obviously, some of the things you would have seen is domain names, IP addresses, and so forth.

Now I'm going to walk through the core IANA function areas, and speak to them in a little more detail. This started as an informal grouping, but it has certainly become formalized over time. We divide them into three functional areas. We have protocol parameters, number resources and domain names.

I'm going to start with protocol parameters, firstly because essentially everything on this slide is a protocol parameter, it's just that number resources and domain names are specialized applications of protocol parameters. Protocol parameters are essentially used in all kinds of Internet software. Really, every device on the Internet is using them in one form or another.

Protocol parameters today are issues directly by IANA. Those that wish to use them would approach us with a request directly, and we would issue them. Rules differ depending on what kind of unique identifiers we're talking about. Applications are evaluated by our IANA team. We do it according to set criteria. Generally speaking, the creator of the unique identifier will specify the requirements for who is eligible, and under what terms.

I mentioned earlier that protocol parameters visibility is limited, generally speaking. Software implementers would be the only ones

that see it, typically residing only in software code. Just to give you a sense of some of the things you might not have heard of that are critical parameters on the slide, IP header flags, port numbers, media types.

To explain what some of these are, port numbers is one of the primary mechanisms by which two devices talking on the Internet discern between what is, let's say, a web page being transmitted, versus an e-mail being transmitted, versus a telephone call. Port numbers are one of the essential ways that that kind of traffic is discerned in a transmission.

Media types is how, in that transition, file types are negotiated. If you've ever wondered how your web browser knows that it's receiving an image versus a sound file, versus a movie, versus a document, media types is how that is done. We have a registry of all the different filing codings, and that is the media types registry. We have a registry full of port numbers that assign port 80 to web transmission and port 25 to e-mail transmission, and so forth.

Where do these numbers originate? Well, it's predominantly through Internet standardization. The preeminent organization for standardization for the Internet is the Internet Engineering Task Force, more commonly known as IETF. They create the standards documents that drive unique identifier creation. Generally speaking, these documents have a section in them called IANA considerations, and that is where the bulk of our work in this area is driven by. It instructs the creation of new protocol parameter registries, and instructs us on the

registration policy. It also specifies initial registration value and reserve values for those different identifier types.

It might be surprising to many, but there are literally thousands of registries. Within that, many more values that are registered. At this time, there are around 3,000 protocol parameter registries for all manner of technologies that are used on the Internet. We have an index on our website where you can scroll through and peruse them if you so desire.

Our role, fundamentally, when it comes to protocol parameters, is maintaining those registries, which includes keeping the registry data accurately and publishing it on the Internet for others to use. Receiving and evaluating requests from different parties who wish to have new registries created, or have modifications made to existing registries. For example, adding new values, or perhaps modifying purposes. Also, we provide advice back to the IETF on their upcoming standardization efforts. Wherever there is an upcoming Internet standard in development in the IETF, we're consulted. We provide advice back on how it might be implemented as part of the IANA functions.

That's protocol parameters. Now, I'm going to move on to number resources, briefly. As I mentioned before, number resources and indeed domain names as well are specialized forms of protocol parameters. In the case of number resources it really falls into two different groups, here. One is IP addresses. IP addresses are the unique identifiers for devices that are connected to the Internet. Autonomous system numbers are unique identifiers that group networks on the Internet.

The way I like to explain it is you can consider AS numbers as postcodes or zip codes that aggregate groups of IP addresses into single units.

Number resources are predominantly allocated hierarchically through a system where we work closely with the five regional Internet registries. Essentially, what happens here is we allocate large blocks of IP addresses to regional Internet registries. They in turn allocate smaller blocks to ISPs and network operators within their region. Those ISP and network providers allocate individual IP addresses to their customers. Some specialized allocations we do directly, but these are not the kinds of allocations that you would regularly use. These are highly specialized applications, Multicast being one.

What I think is particularly unique about number resources allocation, as opposed to the other things we do, decision making is highly deterministic. There are formulas and qualification criteria in all the global policies that inform this line of work. A recent change here that we implemented in the last few weeks was a new RIR dashboard that provides public explanation of eligibility requirements, and is updated daily to show which RIRs are eligible for what kinds of allocations.

Finally, moving on to domain names. The most famous IANA function, and the one that probably most people here are familiar with, is managing the DNS root zone. DNS root zone is the top of the DNS hierarchy. It is the official record of which top-level domains have been allocated. It is also the database upon which DNS queries actually function correctly, by directing DNS traffic to the servers of TLD operators.

Like number resources, these are hierarchically allocated. IANA is responsible for the uppermost level of allocations. Probably very familiar with the diagram of the root tree, but TLDs sit under there, second-level, and so forth.

What are the IANA tasks when it comes to domain names? The final way, the most critical function we have here, is receiving and evaluating root zone change requests. We receive them predominantly from TLD operators. We evaluate them against the policies that are in existence as well as our operational requirements.

Some of the types of request we would receive of this nature. Firstly, the original assignment of TLD. If someone wants to create a new TLD, that request is sent to us and we evaluate it. If you wish to transfer a TLD. This is when one organization wishes to hand operation of a TLD to another. That is a request that we would evaluate. There's routine maintenance of the technical infrastructure for a TLD. It might need occasional changes to the name servers, and other technical elements.

We also maintain a set of points of contact. These are the parties that we recognize as being able to act on behalf of a TLD. Whenever those points of contact need to be changed, that would be coordinated through us. We evaluate requests and those that need to be implemented in the root zone file itself. We transfer implementation in that file, and then onward to the root servers.

Some of the other domain names functions not directly related to the root zone include .int. .Int is a very small registry of around 200

registrations. Today, it is solely limited to inter-governmental treaty organizations. This is a registry that has strict eligibility requirements, and something we continue to maintain.

We also maintain what is today called an LGR repository. You might know it better as an IDN table repository. Essentially, this is the policy upon which TLD operators allow IDNs to be registered in their zones. It essentially documents which characters or code points are allowed for a given language or script that is supported by a registry.

One other, I think, fascinating area of operation that is relatively new to IANA ... It's actually 10 years old next year. Relative to the other things we do, it's relatively new. It's maintaining the root zone key-signing key. To translate that, we typically refer to it as the trust anchor. This is the key that enables DNSSEC to successfully work. This key is special because it needs to be maintained in a very secure and practiced way. It is very hard to change. All the other keys that are using DNSSEC can be relatively easily changed by communicating the new key to a business partner.

Because this key is at the top of the hierarchy it is typically hard-coded in software and is essentially unchangeable without a very carefully orchestrated process. Therefore, it is very essential that it be managed in a very specific way. It's also important for trust in the DNSSEC system itself, that it be done in an open, transparent, and trusted manner.

That is why we have a very elaborate process of maintaining this key, that involves public key signing ceremonies. These are typically held

every three months, where we have experts come and oversee operations of the key. Some of these are community volunteers, called trusted community representatives.

Our job here is managing the whole life cycle of this key. This includes whenever we need to create a key, whenever we need to use a key, and whenever we need to replace a key, what we know as a roll-over. These key ceremonies are published online. These days, we stream them on YouTube and interested parties are welcome to attend and observe them.

That is my quick summary of the IANA functions at a fairly high level. Together, these three components, protocol parameters, number resources and domain names comprise what we know of as the IANA functions. Importantly, these divisions also represent the three areas of accountability we have to the community. We're going to get into the details on that in a little bit.

I mentioned at the start PTI, the specially-built non-profit that was created in 2016 to operate the IANA functions. Its role is to perform the functions I just described. It also does things. It hires the IANA staff that performs the role. Structurally, it is a non-profit. It has a single member. The single member is ICANN. The way we typically describe this is it is an affiliate of ICANN.

Here is the IANA staff as of today. We have 15 employees working at PTI, performing all manner of the functions. Predominantly, our team is based in Los Angeles, but some of these faces will be familiar to people

here. We attend different conferences and engagement events to familiarize you with our operations and our team.

There's also a five-member board of directors of PTI. Two of those members are NomCom appointees. This is the current composition of the PTI Board. Lisa Fuhr is our chair. She's a NomCom appointee. We also have Wei Wang, another NomCom appointee. I am one of the three seats of the PTI Board appointed by ICANN. David Conrad is the second. Obviously, we're missing one here. We recently had one of our directors leave the organization. It is on the agenda for Thursday for the ICANN Board of Directors to appoint a new director. That should be the case just for a few more days. Once we have that fifth director, then we'll be back at our full complement.

That's PTI's role. What is ICANN's role in all of this? Fundamentally, ICANN is responsible for the IANA functions. It contracts PTI to perform the IANA functions, and oversees its performance. One of my primary areas of accountability in performing the IANA functions is to make sure we fulfill the contracts we have with ICANN in performing them. It also provides us with a significant number of resources.

We receive shared and dedicated resources from the various different departments of ICANN. The legal services of ICANN, IT in terms of our end-user support, server infrastructure, software support, and so forth. The ICANN HR department manages PTI's HR requirements. Finance, and many other departments of ICANN.

Importantly, ICANN provides all of the funding to PTI. No one else pays PTI directly. All funding sources come from ICANN. ICANN operates some of the accountability mechanisms for PTI. Most notably, this is the Customer Standing Committee. In the naming function contract, the Customer Standing Committee distills the opinions of the naming customers. It makes sure that we're fulfilling the provisions of the naming contract. ICANN also operates the IANA naming function review, of which the first one will kick off in the next few days.

With that, I'm going to hand over to my colleague Naela, who will walk us through some of the details of IANA accountability and performance.

NAELA SARRAS:

Thank you, Kim, and welcome, everyone. I will take you from here to ... Essentially, Kim walked us through the three IANA functions, and defined that these are really the functions that we're accountable for in our day-to-day work. Those accountability mechanisms are defined in different ways, and different documents.

Let me first describe where those accountability mechanisms are derived from. Kim explained what we do in the protocol parameters area, and that we're primarily working there with the IETF community to deliver that service. The Internet Engineering Task Force, here, and ICANN, have a memorandum of understanding that represents a contract, essentially, between ICANN and the IETF. That spells out what it is that the work was supposed to do to perform that function. Then,

there is a sub-contract that is between ICANN and the PTI to perform those functions.

Similarly on the numbers side, there is also an SLA between the five Regional Internet Registries, RIRs, and ICANN, for delivering the numbering service. Again, there is a sub-contract between ICANN and the PTI that basically says, “Deliver those services according to this contract.”

For the naming function, the one in the green, that’s a little different. That contract is between ICANN and the PTI. The oversight of that contract happens through the CSC, and we’ll go a little bit into that, the Customer Standing Committee. That’s the series of contracts that govern our work. Three different contracts that primarily cover the different functional areas that we work on. Next slide, please, Kim.

To stay accountable to those accountability mechanisms, it’s spelled out on all of these contracts that we provide performance reports against those functions. Each month, we produce for the names ... Numbers? Is that where we’re at, Kim? Yes, sorry. We produce reports for the names, numbers, and protocol parameters. We have different sessions or meetings in which we review that performance. We also make those reports directly available on our website. Go to the next one, please?

For all the three functions that we do, we post our websites on IANA.org/performance. If you go there you’ll find all of our reports each month, going back two years. For some of these that have mandated

since the beginning of the transition, of course they go back to 2016. Each of the reports go into the different metrics that we measure against, and how we performed for that month.

The naming community has another mechanism in which they can monitor real-time our performance. That's the SLE dashboard. That's a real-time dashboard that measures the different metrics that we're measured against, upon the completion of each request, and posts it automatically. That exists for the naming function, and the website is available there, on the slide. Essentially, what we provide to the CSE each month is a distillation of what's coming out of here, the bottom chart, which is the dashboard. Next function, please. Next slide.

How do we work with the community to monitor these SLAs? First of all, in the IETF, in the MOU between ICANN and the Internet Engineering Task Force, those KPIs, we're required to measure against. That's what we report on. For the naming community, there are around 70 metrics that we measure throughout the processing of the requests, of different measurement points that we are asked to look after and measure against.

It's around 70 where it stands now, but we should note that the Customer Standing Committee is very active in looking at if these are the right measurements, should more measurements be added, should some measurements go away if they decide that it is no longer measuring meaningful information?

One of the areas that was clearly missed, I think, during the transition talks, is there were no measurements against the IANA processing of LGR tables, or IDN tables, which Kim talked about earlier. That's in the process of being added as a metric to measure the IANA performance against. That number will go up.

For IP addresses, also in the SLA between the five Regional Internet Registries and ICANN, there are specific measurements that we perform against and report on. Each community sets up its own mechanism for how they want to look over the performance reports that we provide and hold us accountable to that. The Customer Standing Committee meets every month. They go over the performance report and we explain the performance for that month and any misses that happened during that month. Then, the Customer Standing Committee in turn produces its own report reflecting on the report that it received from the IANA function. Those are all available on the SCS website, as well, their own reports.

The IETF meets regularly with the IETF liaison, that's an IANA staff member, and go over the performance as well. RIRs have set up something called the IANA Review Committee. That's a 15-member committee. They receive our report monthly, and I think annually they sit and review the reports, and then they issue their own report, which they also post on the NRO website. All that performance is further monitored and reported back to us as whether or not it's keeping up with the expectations of the contracts. Next slide, thanks.

In terms of how our performance is, and in the actually day-to-day, I run the team that does all these requests that come into IANA. We're the ones that are receiving new requests and processing them. Starting here with the protocol parameters area, it used to be that we internally looked at either the work that comes from the IETF, or protocol parameters, as they are our most voluminous area. We used to say if you just strictly go by number of requests, it was the IETF work that represented the highest numbers up until now. As I will show in the next slide, that has been taken over by the naming function.

With IETF requests, the numbers that you see here is pretty much how the performance has been for years. It has stabilized. It has a little bit [specks] here and there, but it's pretty much the number that you can expect in a given month, in the low hundreds, high 90s.

At the end of each month we distill all that data and we provide the reports. Each month, we measure ourselves against the KPIs that are defined in the MOU between ICANN and the IETF. We should also say that we have an internal KPI that has set a little bit of a higher bar, so that we can keep ourselves in check. If we start missing our own internal SLA, then we know something is going on, and we need to pay attention before we start missing the bigger SLA, or the one that we report on.

Again, we're using the same check here to indicate that it's a satisfactory performance, that our customer base, in this case for example the IETF, is happy with the performance, and we are continuing to perform according to the KPIs. Next slide, please.

Then, we come to the domain names area. Here, I said it used to be protocol parameters that was, in terms of volume, our highest. Now, it has easily been taken over by domain names. Especially in the last few months, I would say that has been the case. This is natural. The root zone has grown drastically in the last three or four years. We've gone from around 300 TLDs, around 2012-2013, to now we have 1,500, give and take a few. We have 1,500 top-level domains.

If you remember back to what Kim said, in terms of our role in managing the root zone, we have the strict requests to actually put the TLDs in the root zone. Then, there's the maintenance that happens each year on those top-level domains, be it name server changes, contact changes, just merely doing the KSK rollover each year. That generates a number of requests for us each year.

This one anomaly here, in August 2019, we had 1,017 requests, with a little lightning bolt going through it here. This represents a contractual requirement that ICANN had for gTLDs specifically earlier this year, which was RDAP server. All Generic Top-Level Domains are required to enable and serve data via RDAP, now. In August 26th, 2019, all gTLDs were required to add an RDAP server link to their IANA RDAP registry. As you can imagine, all these gTLDs that are now sort of around 1,200 TLDs, sent us requests to put their links in the IANA registry.

This is a good chance to say and to shamelessly plug in that if you have such a type of request, if you are a large registry operator, and you come to a point where you need to do a roll-over of your DNSSEC keys, or you need to do a bulk-change for all of your TLDs, please talk to us. Please

coordinate with us ahead of time, if possible. August was a tough month. People go on holidays, people have kids out of school that they want to be home with. In anticipation for August, we had to ask people to stagger their holidays and be available, because we knew all of these were coming to us in August.

This reflects the domain names area. This is strictly taking data out of the TLDs. The same thing had to happen for registrars as well, and those numbers are now reflected, here. Registrars also had to list their RDAP links, or have to insert RDAP links in the Registrar ID Registry. Again, we've had to pretty much figure out new processes for receiving those links from customers and putting them in the registry and attesting them, and reporting back to them when they didn't pass the test. Quite a bit of operational work, there.

In the SLAs [mat], here, we put strictly the number that we report in the report to the CSC about the percentage of SLAs that we met, of those around 70 that we measure. The checkmarks mean for that month we received either a satisfactory or even an excellent. For 100% the CSC grants an excellent performance rating, and anything below 100 is satisfactory. If it's grossly not satisfactory, they have a different rating, which I believe is "needs improvement," and hopefully we won't. We have not reached that level.

That's it, I think, for the names. Let's see. On the numbers side, the volume is drastically smaller. In fact in terms of operation work, this is our smallest area, for me and my team. As Kim said earlier, in the presentation, the allocations here that happen right now are in V6 and

AS allocations. When they allocate a number, it's quite a large number that takes a while for them to go through and in turn allocate to their customers. It takes a while for them to come back.

As of March 2019, we had a small pool of IPV4 addresses that we were still allocating on a schedule basis. That pool is now down to a level that we can no longer allocate equally to the five Regional Internet Registries, so that pool is empty. We only allocate ASN and IPV6 addresses. We've had a little bit of the IPV6 addresses pick up while we allocated them /12s to some of the RIRs. It's fair to say that it was a new process that I've done for the first time this year. I think the same for Kim. We've been with ICANN for a long time and PTI, so it's something that happens very rarely because of the nature of the block that's allocated.

Again, the performance the SLAs met is according to the agreement that the RIRs and ICANN have, and in the sub-contract with PTI, and that the checkmark is satisfactory in that we met all the SLAs. The IANA Review Committee has put at least two of the reports on the NRO website, that are available on the NRO website. In their findings, the PTI has been performing the IANA function, the numbering function satisfactorily, and there have been no issues.

With that, I will turn it over to my colleague Maurilia, who will take you a little bit more through other efforts in which we do continuous improvement in addition to monitoring our performance and reporting on it. Maurilia?

MAURILIA GOMES:

Thank you. Hello, everyone. Just to piggy-back on everything Naela was mentioning, we are accountable to our contracts. Therefore, we always need to be monitoring how we perform our work. My job is to make sure that we have these programs in place and that we have controls in place to ensure that our processes are efficient, our systems are secure, available, and we operate with processing integrity.

How do we do that? We have audits in place, customer satisfaction surveys. We have a structured process going on across IANA. Of course, we plan. We have a culture of continuous improvement within the team. I'll go over a little bit more in detail how those programs work. Can you move to ...?

We have two information security audit programs. We use the SOC 2 and the SOC 3 framework for those. It's led by a third party firm. The SOC 2 audit is a comprehensive report of the systems that we use to process the requests, that Naela was just talking about. The SOC 3 is a public report. You can access it on our website. We had it issued without acceptance since 2010, when we first started conducting this audit. It audits the security control that governs the root zone key-signing key.

Those two programs help us improve, especially the SOC 2. The audit firm comes to us and tests our controls, as far as security, availability and process integrity of our systems. We use that report to help us stay on top of our systems.

We also measure customer satisfaction in two different ways. We have an annual survey which is administered by a third-party vendor. We've been doing that since 2013. This year we switched the approach to focus on engagement with the community. This year, we got a response rate of 3%. The overall satisfaction, on a scale from 1-5 was 3.6. We will be publishing the report in December, so stay tuned for that if you have received the invitation to participate and you're interested. We publish it on IANA.org as well.

The other measuring indicator that we have is through the "how did we do" survey, which we launched earlier this year. We send out a survey after a request is resolved. It's just a simple, "I had a good experience," or, "I had problems," is one question that, once your request is resolved, you will receive. We address the feedback or escalate it within a matter of days. It was a switch on ... We used to do this annually. Based on feedback from you guys, we changed it to be more time-bound. We can work on those improvements more immediately than waiting a whole year to address it. So far, the average monthly satisfaction rate on that was 86%. The response rate is 36%. We get a lot of responses on this survey.

We're still working on some improvements on the tool because a lot of the dissatisfaction that we get seems to be around dissatisfaction with a policy that we cannot really control, versus the service that we delivered. We're working around the tool to see how we can separate these two and get a more accurate satisfaction rate. There's work being done with the tool, but that is a program that we use to drive improvements as well.

This is just an illustration on the survey results. This is the “how did we do” survey. We break it down. We have all segments, and we have broken down by domain protocol numbers. The questions one, which is not here, is just general questions we get on everything. You can see that the satisfaction is pretty high there on all of the functions. This is for the past three months. It’s something that when we do get a dissatisfaction and we get a response from a customer, we immediately go in and reach out to the person who responded and try to address it right there and then. We don’t like to wait for it. If you have a request and it gets resolved, you get a survey. It’s one question, it doesn't take long to respond to it.

This is just to summarize. We do focus a lot on monitoring the work that we do to make sure, besides just the SLAs that we’ve had with the community since 2007, that we are working with the most efficient systems, that we have documented processes in place, and that we’re constantly reviewing it based on the ever-changing needs of the community.

We do use the EFQM model, if you’re familiar with it. It’s a quality management framework that we received a Committed to Excellence award in 2013. We use that to assess our department from processes, to system, to leadership. We assess the whole department using the model. We use the results of that assessment to drive our improvements as well.

The survey we’ve been doing since 2012, as I mentioned. All of the input that we get from these audits, assessments, and surveys, we use it to

drive improvements to the systems, processes, and tools that support our work, which Kim will go over next. The systems like the ticketing system, the RZMS, the Root Zone Management System, are a further automation that we get input from the customers that they would like to see. We're also working on a new protocol parameter management system. This is all work that stems from the input that we get from the community. We're very serious about monitoring and improving the processes that we have so far. With that, Kim will be talking more about that prioritization and planning.

KIM DAVIES:

Thanks, Maurilia. Okay. Moving on to some of the forward-looking items, like how we prioritize, plan and so forth. The first thing I wanted to bring your attention to is strategic planning. PTI has in its bylaws a requirement, in fact, that there is a unique strategic plan for PTI as it pertains to delivering the IANA functions.

Now, for the first three years since the transition we already had a template of what we needed to accomplish. It was the transition plan. The transition plan set out a number of milestones over the first three years that needed to be accomplished. That has been the basic operating strategy for PTI. Now that we're at the three-year mark, and we've accomplished everything in that document, we are now looking at what is the longer-term strategy for PTI.

The PTI Board has taken this on this year as its project. It is now leading an effort to develop that strategy document. We've already identified

some of the key objectives that we want to reflect in this document. I'll just briefly walk through them. Obviously, we want to continue to focus on what our customers need from the service, that is key. You just heard from Maurilia some of the culture that we have in our team, in terms of continuous improvement, operational excellence, and so forth. We want to maintain that and have programs to support that.

We want to maintain the community's trust in our delivery of the service. We want to add value and usability. If we're not adding value then it's not useful that we continue to do things. We want to have an ongoing focus on the evolving security requirements. As with any service provider security is an essential topic that you need to not be complacent with but continue to monitor and adapt to. We need to have support for that.

Also, one thing that emerged, and this is kind of a meta-process, PTI does have a need for a separate strategic plan from ICANN. There are benefits in aligning it with ICANN's strategic plan. Firstly, we are a subset of ICANN's responsibility. Clearly, we must fit in within ICANN. Secondly, the community gets exhausted by all the discussions on operational planning, strategic planning, and so forth. The more we can align those processes to deal with them all at the same time through the same processes, the better for everyone. We're going to take this opportunity to work out how we can best align what PTI requires in its bylaws with what ICANN requires from its.

Also, another thing we identified is we don't really have a clear vision statement for what PTI is trying to accomplish. We do have one that was

crafted by staff many years ago on our website, but this is actually an ideal opportunity to reevaluate that and try to come up with a crisp definition of what the community actually sees IANA is here for, and striving to accomplish.

Some of the key milestones to be mindful of, our goal here is to share a draft of the strategic plan by March next year. Just in time for the next ICANN meeting. We're aiming for it to be effective next fiscal year, starting July 1st 2020. Obviously, if we receive feedback after March that suggests we're quite off-base, we'll take it back for further drafting and review. If it is well-aligned with what the community expects, then we will put it into effect.

Budget development. PTI also has its own budget. It is a sub-set of the ICANN budget, but nonetheless we go through also a separate process for developing the budget. There are actually two budgets. One is the PTI budget. This is the essential cost for delivering the IANA functions. Then, there is an IANA budget. What the IANA budget is is the PTI cost, plus what it costs ICANN to enable the IANA functions, things like how much it costs ICANN to oversee PTI, to manage the contracts, to perform the community reviews that you heard earlier, and so forth. We do budgeting very early in the cycle. We're required by the bylaws to have a draft budget ready nine months before the start of the fiscal year, which means that September 30th we need to have it ready for the next fiscal year, starting July 1st the following year.

As of now we've gone through that process for fiscal year 2021. Our budget is currently out there for public comment. I encourage you to

contribute if you wish. The draft budget that we've put for public comment is based on the priorities discussion that we had during the summer time. The high-level summary here, and I don't want to give short-shrift to the budget ... We do give detailed presentations, but this year we're appraising a budget that is roughly consistent with last year. We're not undertaking any ambitious new spending. The headcount is stable. We're really looking at doing something very similar to the previous fiscal year.

Now I just want to talk about some of the projects we have ongoing with the team to evolve and improve our services. Generally speaking, we can divide development activity into three different areas. Firstly, we have the technical projects that we run. These are tools and system-enhancements that you heard a little bit about before. These are led by an internal development team we have of two people within the IANA department.

We have shared projects with ICANN's engineering and IT department. These are on more of the bigger systems that are often shared with other parts of the company. We have continuous improvement that we do on key management facilities, and the key ceremonies. We have a two-person Cryptographic Business Operations Team there that does that.

Then, on Naela's side of the fence, we have operational projects. Here, we're looking at continuously reviewing and refining how we do our core business, how we process those change requests you heard about earlier. Also, implementing the outcomes of those audits and reviews

that are conducted on a regular basis. When we have a finding that there is an area that warrants examination or further improvement, we'll take that on and try and work out an action plan to make it happen.

Then, we have other strategic projects that affect the IANA services. Just to give you a sense of some of the things we have on the boil right now, I'm going to go through them briefly. The ones in bold, here, I'll go into a little more detail. You heard earlier we just launched a new Internet number resource dashboard. Future KSK rollover planning. We're making improvements to how we handle smart cards in the key management facilities, as well as upgrades to the safes that contain the key itself.

We're building a next-generation root zone management system right now. We're developing new key management for the KSK. We're designing a new authorization model for the root zone. We've built an RDAP server for the IANA resources. This is TLDs, .int domains, and the number allocations we made to the RIRs.

We are working to improve that “how did we do” tool to fix some of the issues we've seen in the early deployment. We're currently in the process of replacing some of the trusted community representatives that we have at the key ceremonies.

We've recently done a .int zone inventory. We review every single .int domain, make sure they're still valid, and taking steps to improve that zone. Platform upgrades for reverse DNS. This is something that applies

to Regional Internet Registries. We're building a new registry workflow system for protocol parameters. I'll talk about that in a moment.

Improving the internal tooling, the stuff that our staff use day-to-day to do their work. We run a certificate authority within ICANN for a number of things, like the Trademark Clearing House, and a few other things. We're working on re-planning how that's configured. We'd like to evolve the IANA website to make it a bit more modern and responsive to customer needs. That's something we are working on.

We're working on key management facility monitoring so that we have more transparency in when it's accessed, when we have security alarm events, and so forth. We're improving the ccTLD transfer process. This is a process very few people use, but those that use it know it's a very document-heavy process. Some of you would like to use modern document management tools to do this, as well as creating new checklists and form-based approach to receiving the information.

Strategic plan development, I already touched on. Next generation protocol parameter reporting of the three kinds of monthly reports we do. The protocol parameter reporting is the oldest. We can take the learnings we've had from some of the more recent reporting, like we do for the CSC, and try and apply those learnings we've had to how we do protocol parameter reporting. We're also trying to improve our internal QA process, how we have staff check on others' work to make sure that areas don't slip out of our processing chain.

Root zone download service, today. The root zone is published on the root servers, but it's also available for download. Today, that service is provided by Verisign, but we're migrating that in-house to be operated by IANA. We're improving the operator request process. This is the process by which root server operators communicate their needs to IANA.

In brief, we have a lot of things going. This is just the stuff that I would consider active. We obviously have a lot of things we would like to do that we haven't quite embarked on yet.

I wanted to go into a little bit more detail on just a few of those projects, because I think firstly they're timely or they're particularly critical to the delivery of the other functions.

The first one I wanted to talk to was KSK rollovers. For those that are not familiar, the first KSK rollover project concluded recently. In August 2019 we finally deleted the original KSK. This marked the culmination of our ... Gee, it's not even five. Maybe a seven-year project, if you counted from beginning to end. The predominant day, the most important day, was October 11th of last year. That was the date that the KSK was literally cut over from the very first one to the new KSK. The project was widely considered as a success. Certainly, there was very minimal disruption to Internet operation on that day. It was well-planned, well-communication. We had a lot of involvement from the technical and operational communities to make it a success.

Based on that experience, we're now thinking about how to do it another time, and another, and another. What we're really trying to do is normalize operations here to create a consistent approach, ideally with a regular cadence. What we've proposed is we put out a paper just a few days ago for public comment. It spells out the details on how we plan to do this. Some of the elements of this is a regular cadence. We propose a regular period in which we'll change the KSK over and over again.

We're also increasing the amount of time we have a key in standby state. This gives us better resiliency to an unanticipated event. It's never happened, but should there be a need to do an unplanned change to the KSK ... Let's say there's some kind of emergency, we'll be in a much better posture to do that if we've already created the key in the first place.

There are some other tweaks to the process, but in general, because the last rollover was quite successful, we're looking to not change too much and use that as a model moving forward. We put that for public comment on the 1st November, and it's open for the next couple of months. We encourage those interested in that to provide comment. We'll give a more detailed presentation of the details of this to the DNSSEC workshop later in the week.

Root zone management system. We've had an automation system in place for roughly 10 years now, that allows TLD managers to perform self-service, lodge chain requests for their TLD that then goes to staff for processing. We continue to make updates to that system that is now

10 years old. The latest release was in July, where we added additional RDAP service improvements.

However, we realized a couple of years ago that the current system was getting to a point where it was no longer fit for purpose, in that we wanted to make some radical changes to how we deliver service that the technological choices in the old system couldn't easily accommodate. A decision was made to build a new service essentially from scratch. That work is well underway. I would say we're probably 75% of the way there.

The baseline there is obviously to keep all the existing functionality that TLDs rely upon, but also to add new functionality. Some of the initial functionality we're putting in the system for its first release is a new user model. This allows the creation of user accounts in the system. Two-factor authentication and other security improvements is one aspect of that user model. Another aspect that's not on the slide is having better "know your customer" procedures, so that we know the person behind an account. If a credential is ever lost, we're able to reestablish trust, reissue credentials and so forth, in a safe way.

Another feature is we're separating out the technical check component of the system. This is where we check the servers and the operations of a TLD manager to make sure they're working correctly. Our experience over the last 10 years is this is the part of the system that most needs changing on a regular basis, based on developments in the industry and so forth. Rather than having that fully integrated into the core system, having it as a separate service will allow us to be much more responsive

to changes, and be able to deploy them much quicker. That's one of the architectural changes to that system.

We're adding better support for bulk updates in the system. If you think back 10 years, Naela mentioned back then we had about 300 TLDs. Generally speaking, there were about 300 TLDs and there were about 300 companies that ran them. Any individual company typically ran one TLD and that was it. Sometimes they ran a couple, but most companies ran one TLD.

The model today is very different. You have some vendors in this industry that run hundreds of TLDs. Our system isn't well equipped to cater for them. We really had a mindset when we built the original system that any one company is running one TLD. A lot of the interfaces are designed around that assumption. With the advent of these large companies with a portfolio of many TLDs, we're building better tools to support bulk updates.

The first step in this journey is to add an API. For those not familiar with that, that allows programmatic access to the system. This will allow those companies to write automated software to talk to our system, to make automated updates to the TLD. We're also at this process going through and refreshing the user interface to the system, adding things that weren't really critical 10 years ago, things like mobile phone support, and so forth.

Another thing that this rebuilding of the root zone management system has triggered, but it's not exactly tied together, is the authorization

model. We have an authorization model that's kind of rooted in the 1980s, today. Every domain has an administrative and a technical contact. Those two contacts are required to approve root zone changes for that TLD. We're migrating to a process where we have a much more flexible model, where we have users of the system, it can be one, two, three, four, 10, 20, that can operate on the TLD.

We want you to be able to configure levels of access. You might give a certain staff member full access to manage the TLD. Other staff members have limited access. Maybe you give limited access to a vendor to do a certain change request. These are the kinds of operations that we want to be able to accommodate. Currently, the system is like one-size-fits-all. The new system should be a lot more flexible, and that's the goal here.

We're also changing some of the business logic. Not a lot. One key area I wanted to highlight is today, if a name server is shared by many different TLDs, every single one of the TLDs that uses it has to approve a change to it. We're looking to change that model to make it a little more streamlined. We're going to be talking about that in more detail to the ccNSO and GNSO later this week.

TCR replacement. Those trusted community representatives that observe the key ceremonies, we're looking to replace some of them partly due to retirements of the volunteers that do that. You might know of these people as the seven key holders. Usually when it's characterized in the media they're referred to as the seven keys to the Internet. We call them trusted community representatives. Most of

them have served since 2010. Given we're coming up to the 10-year anniversary of the KSK, many are considering this time for them to move on to other things. We need to replace them.

We've launched a new evergreen process where anyone is welcome to volunteer to be a TCR. Currently, over 100 people have already volunteered, which is great. We have objective selection criteria that we try to apply. We score candidates in terms of their diversity in geography, culture, skills, experience. We also assess them based on their reputation and standing within the community. Fundamentally these people are meant to represent the broader community to ensure they trust the process. They act as their agents. We want people that are trusted in the community.

Just in the last week we finalized our first selections under this new process. They've been advised of their success. As we replace existing TCRs, we'll continue to appoint more TCRs via this mechanism.

I mentioned that we have a root zone management system that supports our most critical domain names function that we operate. We do not have a comparable system today for those 3,000 registries we maintain on behalf of the IETF. We're currently in the midst of a multi-year project to build such a system. We're taking learnings we got from the root zone management system. We're also taking learnings from a system the IETF operates called the data tracker. A lot of the work to date that's gone into this system is normalizing all the data that we have. We have data spanning decades of registrations.

Formatting of registries was not consistent. Over time, protocol authors have been creative in the way they record values. We've entertained that and supported that over time. This project has triggered us to normalize the registries, to put them all in a common, standardized format to allow for us to build systems around them.

Our team has valiantly done a massive data harmonization effort and inventory of the data that we had to make that possible. Now, we're in a good place. It is all in a standardized, structural system. We're building the user interface around that and the workflow components to act on that data.

Our aim here is to launch a system that supports a subset of those protocol parameter registries. In fact, we're proposing to start with just one, initially. Next year, we plan to use this new platform for private enterprise numbers. You probably haven't heard of them, but it is our busiest registry. It's a high-volume registry, but it is low complexity and allows us to test the system out before we move other registry types to the system.

Now, to summarize this presentation, and to reflect back on the question I posed at the beginning. What's changed, and what's stayed the same? I'll start with the changes. What's changed in IANA operations in the last three years versus the period before that? I think it's fair to say we have more direct accountability to the community. Prior to that, when we were contracted by the US Government, our contractor was the US Government. We were two steps away from the community.

Now, we provide service directly to the community without that middle step involved.

We've managed to streamline processing of root zone change requests. Previously, the US Government was required to authorize every change to the root zone. That's obviously no longer the case. That allows us, once we believe a change is in compliance with policy, to directly implement that. That has allowed us to streamline the process. We now have SLAs across all three areas of our operation that are mutually agreed with those communities. Whilst we did have SLAs with the US Government, they were accountable to them. In fact, our SLA reports were confidential. They weren't seen by the community. Now, all of our reporting to the SLAs is in public documents, as you saw earlier.

We have a separate legal entity. We have a separately defined budget, as you heard. We have a board of directors. We have other governance mechanisms in place that we walked through earlier. There's more public accountability and accounting of our performance. We have both monthly and in some cases real-time reporting. Today, we have the IANA staff directly employed by PTI, whereas in the past we were employed by ICANN.

They're the changes that we've seen. You'll probably agree the changes that have happened are for the better. Hopefully, not for the worse. What's the same? Well, it's the same staff team. Obviously, we've had a little bit of turnover, but we have staff that have worked there for 10-plus years. We still sit in the same office we did before. We're still providing the same functions. There's been no scope changes to the

functions, and it's still the same customers out there that we're servicing.

We're continuing to work closely with our colleagues in other ICANN departments. We do rely upon ICANN's resources to do our work. Just to pick on one example, our regional engagement, global stakeholder engagement team at ICANN helps us with that. We've had no adverse changes to our core request processing and customer experience. I think any changes there have been for the better, not for the worse. The scope of the IANA functions hasn't changed. I think the transition as trying to preserve the IANA doing what it had been doing historically, and that has been the case.

My final slide is, what in practice has the transition meant? IANA functions continue to be provided dependably to the community. I think our experience in terms of our performance rating, adherence to SLAs, customer feedback, and our annual surveys all reinforces this. The transition process has tailored many aspects of the governance to what the community wants, in terms of accountability of the IANA functions. The IANA, as we've seen in the surveys, continues to rate with high levels of satisfaction and high levels of adherence to the SLAs that have been defined by the community.

That's the presentation. Thank you very much for following along. We have some links on the slide here that will take you directly to some of the elements we've been talking about. We have 15 minutes left. We're very open to having any questions that you might have. Feel free to bring them to our attention. I know we have questions from remote

participants, so while those in the room collect their thoughts about questions they might want to ask, we will have some of the remote questions. I think out ... Yes.

UNIDENTIFIED FEMALE: Here. We have two questions in the remote participation. The first one, Kim, is, “When will 100% transition of IPV4 to IPV6 ... When is it expected?”

KIM DAVIES: Wow. I'm not much of a philosopher or a ... I have an opinion. I'm not sure my opinion is particularly relevant. The truth is that IPV4 is not being replaced by IPV6 but augmented by IPV6. I think when it comes to IANA the important message to share with you is that we have no IPV4 left. We've effectively allocated every IPV4 address we've had to RIRs. In turn, some of the RIRs have run out of their allocations and the rest have dwindling resources. As to when IPV6 will fully replace IPV4, I wouldn't hazard a guess on that question.

UNIDENTIFIED FEMALE: The next question is, “What is exactly IANA/PTI's role in the re-delegation of a ccTLD? Secondly, in case of non-agreement of two parties, can IANA or PTI perform an intermediate role between the parties to settle the re-delegation issue?”

KIM DAVIES:

For ccTLD delegations and transfers, we formerly called them re-delegations. We now call them transfers. Our role is to make sure that the request firstly abides by a consensus that's been reached in the country. Secondly, that the proposed manager of the ccTLD meets certain baseline skills to run a TLD. Predominantly, the due diligence is around consent and consensus in the community. In the case of a transfer, we're looking for a consensual transfer as required by the policy, and consensus of significantly interested parties in the country to such a change. Based on the second part of the question, generally speaking our job is not to be a decision-maker. Our job is to reflect that decision making has happened in the country.

In the case where there is a contention in the country, that there is no agreement, if you look at RFC5091 and other documents, fundamentally our role is to drive that question back to the country to resolve within the country. We're not here to make judgement and make a call from one party to another. What we can do as a neutral party is help facilitate that dialog, if we can. If there's areas of contention where having IANA as a neutral party involved can help address, we're happy to do so.

We can also help by explaining our criteria, our rules, and so forth. That can help move the conversation along, as well. We can also share our experience from other ccTLDs where appropriate. What other countries have done in the past can be instructive for the scenario. We can help share our experiences, there. Generally speaking we're here as a facilitator, but not as a decision maker.

UNIDENTIFIED FEMALE: Any other questions in the room? If you are sitting in the chairs over there and you want to ask a question, if you could move to the desk and ask on the desk microphone, that'd be great. Thank you.

KIM DAVIES: Well, we were either very comprehensive or this is the last session of the day and everyone's ready to go out.

UNIDENTIFIED FEMALE: To the circus thing.

KIM DAVIES: Thank you all for sitting through the presentation. It has been our pleasure presenting it. If you have any questions later on, feel free to approach us in the corridors or shoot us an e-mail. We're very happy to answer them. Thank you.

UNIDENTIFIED FEMALE: Thank you.

[END OF TRANSCRIPTION]