
ICANN85 Mumbai | CF – ccNSO: IANA Public Records Study Group Work Session
Saturday, March 7, 2026 – 16:30 to 17:30 IST

CLAUDIA RUIZ

IANA Public Record Study Group Work Session. My name is Claudia Ruiz, and I, along with my colleague, Joke Bracken, are the participation managers for this session. Please note that this session is being recorded and is governed by the ICANN Community Participant Code of Conduct, the ICANN Expected Standards of Behavior, and the ICANN Community Anti-Harassment Policy.

Please observe the following guidelines to participate in this session. I will also post them in the chat for your reference. During this session, questions or comments submitted in chat will be read aloud if put in the proper form, as noted in the chat. If you would like to speak during the session, please raise your hand in Zoom. When called upon, virtual participants will unmute their microphone. On-site participants will use a physical microphone to speak and should leave their Zoom microphones disconnected. For the benefit of other participants, please state your name for the record and speak at a reasonable pace. Thank you. And with that, I will now hand the floor over to Jordan Carter, Chair of the study group. Thank you.

JORDAN CARTER

Good afternoon, everyone. My name is Jordan Carter, .au, Vice-Chair of the Council and Chair of this study group. It's a pleasure to

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.

see you all today. I'm glad you've all chosen to sit in one block to make the most of the collegiality of being able to see each other across the crowded room. Today is the first sort of proper meeting of this group, and I'll just outline what we're going to do.

We're going to take a bit of a reminder about why we're having this study group and what the issue is that has motivated this work. We're going to have a little bit of a look at the project approach we're going to use to do the work of the study group, which will be finished by the end of this calendar year. It's a sort of short-term group, March to December. Because it is about IANA records in relation to ccTLDs, we thought it would be good to begin that fact-finding with a presentation from Kim Davies from IANA about what the records are that IANA keeps with ccTLDs to just ramp up that familiarization that people have.

So that's the essence of today's work, and we will start with just a bit of a reminder about the genesis of this work. I think if you could go to the next slide, Claudia? No, go back. We won't do that one. We'll come to that. That's item four. So a few years ago now, the ccNSO decided to do some analysis of gaps in the global policy framework. That was helped by insights from Kim Davies and the team from IANA who are obviously at the operational forefront of what happens when CCs are working with IANA. They identified some interesting areas where they felt that the policy framework could potentially gain in depth or clarity to improve it essentially and to improve the experience for ccTLDs working with IANA, or to just make the job of administering the root zone for ccTLDs easier.

Could we just flick to the other one of those Google Docs, Claudia?
Thank you.

And this study group approach was chosen because the ccNSO was not convinced that there necessarily needed to be new policy, but felt that it was worth exploring the issues in more depth than a sort of all-in-one quick process. This is an email from some time ago, or a distillation from it on topic areas that have been suggested. In this IANA Public Records Study Group, we're dealing with the first three topics of this document. One is about the accuracy of the records. Are they accurate? Is the data that is published on the website when you go and look for the details of a ccTLD accurate or not? What is the level of accuracy?

This document will be circulated with the slide pack later, so I'm not going to go into huge depth. You don't need to memorize the text or anything. If you could scroll down, Claudia, so we can see all of number one, if that's plausible. So the issue around data accuracy is about to what extent is what is published in the record consistent with reality? There are a few structuring questions there. The second subtopic of the work here, if you could scroll down to number two, is the question of purpose. Why does IANA publish the records that it does? Why does it retain the information it does? What is the use of those records? One obvious use is to be able to identify the organization that is managing the particular domain. But there are a number of questions around the purpose that may guide what is put in as the information in those records.

And then if we scroll down to number three, it's the question of what to do if the data isn't accurate. Is there a reason to be concerned about that accuracy? We think there probably is. What could be done? There's no process really in place for what should be done if data is inaccurate and is known to be inaccurate and isn't being updated or corrected. Those are the three topics. What is the level of accuracy? What is the purpose of those records, and what should be done, if anything, if the data is not accurate? That's just a way to recap the subject matter of this group. That's one of the two sub-documents. There are other issues there; you can read through it at your leisure once we circulate it. The second thing that we thought we'd take a quick look at, and I'm going to hand over to Bart Boswinkel to give us as much of this document as he thinks is necessary, is what, if any, obligations or policy is in place that relates to this at the moment? Bart, do you want to take us on a couple of minute tour of that?

BART BOSWINKEL

Yeah, thanks. Can you go to the other document, Claudia? The second document. This is just done on a basis of a quick search. First of all, what's in RFC 1591, and I think even in Kim's note from almost two years ago, there is a reference to RFC 1591, and you can see the links to its section 3.1. What it states, I will not repeat it. This was refined in the framework of interpretation of 2014. I've even included the link to the document. Can you scroll down? Again, there is a reference to the technical contact and the administrative contact, and there is a little bit in the framework of interpretation

about the ccTLD being permanently in the country. Again, you see the reference. Can you scroll down fairly quickly?

So now another document which is interesting, and I started to think about it was that some of you will know that the ccNSO, one of its first real work items is development of the accountability frameworks. So that's documenting the relation between ICANN and ccTLDs. There are around 55 right now. There is a guideline which is still on the list. You can see the reference, and based on that guideline, there is a template with an accountability framework. There is also an exchange of letters. It captures the same commitments, but then divided over two documents. If you would look at that document, there is a section about commitments of ccTLDs. I will not go into the details of it; you can read it afterwards. Going back in history, then there is a reference to relevant standards. Can you scroll down, please? The relevant standards are captured in RFC 1591.

If you would go back even more in time, and now we're talking about the origins of the ccNSO, if you would look at the scope and what ccNSO is about, you can find it in Annex C of the ICANN bylaws. This was the ultimate conclusion of the negotiations between the ccTLD community or the advisory committee and ICANN at the time about the policy scope at one end, but at the same time about responsibilities and accountability of ICANN and the ccTLDs. If you really want to go into the details, I haven't included it, but go back to some of the presentations from the Shanghai meeting and the Kuala Lumpur meeting. Then you can

see the full picture and how the accountability frameworks and the scope of the policy relate to each other. If you would scroll down, please.

This is purely quotes from the Annex C. Scroll down. This is where we get to. Go back up a little bit. If you look at the ccTLD registry and this is about the policy setting by ccTLDs, one is that the ccTLDs themselves set the policy, or it's happening in country. They have the role of executing it, so really acting upon the policy set locally, and then you have an accountability part related to it. Do you do your job well? And that's partly to ICANN/IANA and partly to the local internet community, including the government. This is where the data accuracy commitments stem from in the accountability frameworks. This goes back to the origins of the ccNSO and the scope in policy setting of the ccNSO itself. That's all. Back to you, Jordan.

JORDAN CARTER

Thanks, Bart. There's a bit of an origin story there, and there's a bit of an explanation given about what the topic is. I just want to take a little pause and ask if there are any questions or comments on that before we move into a look at the project philosophy for the working group's work. I see Nigel Roberts.

NIGEL ROBERTS

Yeah, thanks for that. It's a very useful overview. Some of us predate this and remember how the records were kept prior to the

creation of what we now call the IANA database, around 2000. Without going into detail, it was mixed up with the .com data. What's actually quite interesting is there was a lot of discussion in 99, 2000, 2001 about the meaning of the 1994 requirement that an admin contact be in country.

The conclusion that seemed to be widely accepted at the time was that although John wasn't a lawyer, it was to provide a point of service in the country for a ccTLD. In other words, there must be at least a minimum where somebody could rock up with a process server in American standards and drop a summons or the equivalent on the ccTLD. There's no requirement, absolute requirement, or never was, that the actual organization be incorporated in country, but the important thing is that they have the equivalent of what we call in English terms a branch, or a lawyer or somebody like that who can receive service. It isn't terribly relevant, but as soon as it was mentioned, I thought I'd put that on the record; it might inform the discussion as we go forward.

JORDAN CARTER

Okay. Thanks, Nigel. I think it's a part of the color and flavor of this work. I do not see any hands up in the Zoom room. The first part of this was context. The second part is the project philosophy. The third part is starting to get the presentation on what is actually in these records. So we're going to move to part two. By the way, if this meeting doesn't take the full hour until 05:30 PM, that's okay.

Part two is the project approach that we're going to use. This project approach is designed to deal with an ICANN problem that happens from time to time. Those of you who've been around longer might have witnessed this more times, which is people think they understand the problem they're trying to fix, but as they start trying to fix it, it emerges that they had different problems in mind. There's a project framework called Double Diamond Thinking where you start to get to grips with the subject matter of the issue. You throw ideas at the wall, you understand the context, and it's only once you've begun to do that that you begin to refine down to the pointed questions that you are trying to do.

Then you start to work out how to solve them, and then you refine and settle in on your solutions that you're going to recommend in the findings of the study group in this case. In the slides, there's quite an extensive presentation about how that approach will work, which Bart has prepared. Bart, I'm going to give you a few minutes now to take us on a bit of a highlights tour of those slides. Can you go to the slide deck, please? Think of this as the philosophy of the project with a fair bit of detail. You don't need to memorize it; it will be circulated. We're still to do the work of turning that into a fortnight-by-fofortnight plan for what will be done when, so that'll come up in the next meeting of the working group in a couple of weeks. Can you go to the double diamond? Yeah, there we are. Scroll down.

BART BOSWINKEL

This is what it really is about. If you start thinking about it, study groups are probably the best positioned to do this, and this is something we found out with the other study group as well. The basic idea is that you start really on a discovery path at first, really understanding the problem. It's scoped; there is a scoping document, but there is context to that scope and context to that problem statement as well. Then you start refining it and defining the problem. Then you start coming up with some solutions, and that's the second diamond. Again, you go very broad first, come up with different solutions, and then at the end, you start to deliver. Delivering in this case of the study groups is to start presenting and selecting what you think is the proper way forward. That's why it's called—and this approach has been developed in the context of design thinking. Can we go to the next one?

Applying this approach to what we do right now, again you can see the questions in the discovery phase where we're dealing with facts or fact-finding. This is the fact-finding phase, effectively. Read at it at your leisure. You have some quick key questions for the planning phase. After talking with Jordan, I think it will take about two to three meetings to get this behind us, and then we can start really looking at the data and say, 'What is the real problem we need to address?' That's the next phase. Can you go to the next slide, please?

This is where you start defining. So the purpose is already set, accuracy, and the policy gaps, et cetera. That's something based on the facts you encounter in your journey to address it. Then the

next phase probably will be post-Seville, when you start thinking about solution space. Developing some solutions and then start delivering on it, which is reporting back to the community and to the Council. That's the next phase, Phase Four. Can you go to the next slide, please?

So where we're at right now is really starting and kicking off the exploration of the problem space. That's why you saw these two documents: first of all, the PGA, and secondly, a little bit about the rules and the background of the commitments in the accountability frameworks, et cetera. Back to you.

JORDAN CARTER

Thanks, Bart. So that's the kind of outline of the project. Again, this will be circulated. Have a look at it. You circulated it already? Are you sure? Anyone read it? Okay, good. It has already been circulated, so hopefully some of you—I think, Gordon, you've read it? Yes, this is good. Sorry, that's thrown me off my train of thought. We're going to try a sort of fortnightly cadence from after this meeting to see where we get to on this. If that cadence proves to be too frequent, we can slow it down after Seville.

Really, there's only about three months between now and the Seville meeting, and travel here and travel there and the Easter break takes out a few people. So we'll just do five meetings between the two ICANN meetings to see where we get to in this discussion and moving into the Phase Two work. Then I hope that using the in-person time that we have together in Seville, for those

of us who go to that ICANN meeting, that will allow us to really be very clear about the problems that we're trying to solve. We'll have finished that Phase Two and we'll be starting to work out ways to solve them. I'll just take a pause there and see if there are any questions about that approach. It might be a little bit meta. We'll have a sort of fortnight-by-fortnight plan that gets into more detail about what we're going to do when for the group to review at the next meeting, which will be on the 25th of March. No questions in the room or online. Oh, yes, hi. Just use the microphone, please, and introduce yourself when you speak for the first time.

RAHUL GUPTA

I'm Rahul Gupta. What I'm wondering is, does this only concern the ccTLDs, or is this for registries in general that you'll be studying?

JORDAN CARTER

That's a great question. This is a ccNSO study group, so within the country code domain part of ICANN. Our focus will only be on the records that IANA collects for ccTLDs. When we do things like look at the accuracy levels, we'll only focus on assessing the accuracy of data about ccTLDs.

RAHUL GUPTA

Right. The reason I asked this was because a lot of the results, even though you might focus on ccTLDs, would be applicable to other registries within IANA in general. Some of that is useful because there are issues with management of registries in IANA which will

have a similar flavor of problems to the ones you're going to discover.

JORDAN CARTER

That's an interesting point. Maybe we could grab a quick chat at the end of the meeting so I can flesh out my understanding a bit more of the point you're making, or Kim might be interested as well. But I think one of the things is that as a study group report, it will be published. If there are findings that are of interest or use to other parts of the community, they will be available for them to use. But we don't have the mandate or the scope to be able to work specifically on those other registries or other areas. That's all part of the design of the ICANN system.

Okay. That was part two of three. Part three of three is the enormously exciting presentation that Kim Davies is about to give us, who's just arrived last night from Los Angeles. I think we just now need to move to Kim's slide deck for ccTLD public records. Handing over to Kim. Thank you.

KIM DAVIES

Thanks, Jordan. The intention of this presentation was just to do a lay of the land of what kinds of public records we have associated with ccTLDs and what their purpose or operational principles are. It's very high level; I've probably missed some bits and pieces here and there because I just quickly put this together this morning. Hopefully it's useful nonetheless to stimulate a discussion about

this. It goes without saying we publish a number of public records associated with TLDs, particularly the entity that is recognized to operate a TLD.

Some of the things we publish associated with that are the points of contact that have been nominated by that TLD. We publish objective technical facts, such as what the nameserver records are for a TLD. We also publish user-supplied references, and these are things like the website of a ccTLD, where the registry data directory service is located for a registry, and things like that. We also publish references to any related IANA resources for a ccTLD. If it's a ccTLD that has been delegated or transferred in the last 25 years, there would have been a report that's published, so we link to that as well.

Another thing I would introduce the topic by saying is that we publish the data that we have via multiple different mechanisms. It's not a complete overlap; we publish slightly different sets or subsets of data via different mechanisms. The primary mechanism, though, is our website. We tend to call this the Root Zone Database. We're actually looking at renaming that to be the Top Level Domain Registry, but that's perhaps a topic for another day. We publish directory services two different ways: we have a WHOIS service at [WHOIS.iana.org](https://whois.iana.org) and we have an RDAP service at rdap.iana.org.

When it comes to TLDs, RDAP doesn't actually work via normal discovery mechanisms. I reported to the ccNSO in previous meetings why. We seem to trigger bugs in RDAP clients if we publish

the way that RDAP suggests because we're the root zone, and if the root zone has an RDAP record, it overrides everything else, which is not great. We don't actually have it published in the RDAP directory right now, but if you manually connect to our RDAP server, you can find the information there. Also, we're very close to launching a machine-readable JSON version of the Root Zone Database as well. For those that want to basically compile the information for their own purposes, we emit some PII stuff out of that, but generally speaking, if you want a record of what TLDs exist, there'll be a machine-readable version as well. Sorry, I saw a question in the audience.

JORDAN CARTER

Sorry to interrupt, Kim, but this is interesting. The machine-readable JSON version, are you going to say that is completely free to be used by any third party, such as a ccTLD manager, or even just somebody who wants to publish information about TLDs? .info or something?

KIM DAVIES

I think our general stance when it comes to content on the IANA website is it's in the public domain, and where it's not, we've actually made a statement in the past that you should assume a Creative Commons Zero license for it. Essentially, people are free to do with the data what they like. As a matter of US copyright law,

facts are not considered copyrightable, so we don't have a lot of intellectual property in the data itself, but it's essentially free.

JORDAN CARTER

Thanks, Kim. The reason I asked was I could actually see an opportunity for third parties to add value, given that we know IANA has got a lot to do, and if somebody wants to produce something that gives the data in a much easier all-singing, all-dancing format, it might actually be a ICANNwiki.org.topleveldomain or something.

KIM DAVIES

Something to think about. I'll get to that later in the presentation. Okay, so anyway, that was the introductory slide, and I'm going to get into some of this in a bit more detail as we move forward. Next slide, please.

Let me walk you through individual elements of a typical delegation record that we publish on our website. I'll start by saying that first we simply state that it's a ccTLD, as opposed to other types. Here, we actually historically published more than that. We used to have the country name there, but some fifteen years ago, we removed that. It was essentially because there was some political discussion about what to call certain countries, and to obviate that discussion, a decision was made to remove country names altogether. That's the current state of affairs. Obviously, it's at the expense of some utility that we don't tie what the TLD is

named with its associated country name in the ISO standard, but that's the current state of affairs. Next slide, please.

Then there's the entity name, which is the formal entity that we recognize to be the ccTLD manager. That's pretty straightforward. Next slide, please. Here we have an address. On this, I would observe that, generally speaking, our procedures are such that this should be a physical address as opposed to a service address. We're not looking for a PO box or anything like that; we're looking for a substantive physical location. Coming to the question about presence in country, in terms of our current operational procedures, this should represent a jurisdictional nexus with the country or territory for which the ccTLD is designated. Next slide, please.

Next we have the points of contact, of which there are two. There's an administrative and technical contact. It's been mooted that we should add more. A security contact, in particular, is one that is a regular point of discussion. But as of today, and I think historically has been the case for a very long time, we have these two points of contact. Notably, there is no definition of what these contacts represent. There's no definition of what are administrative matters or technical matters. Essentially, it's in the eye of the beholder as to what the distinction is between these two contacts. Sometimes our customers ask us, 'Well, how should I pick?' I don't really have a good answer to give them. I just say that it will appear in a public record, and it will say it's administrative or technical. There's no further guidance as to what kind of contacts might be made to

these different points of contact or what is expected of these roles. It's just simply not defined. As was noted earlier, there is a policy requirement that the administrative contact must be in the country, so that is something that we validate in the course of our normal operational duties. Next slide, please.

Then there's the question of what is actually in the record itself. It's just a collection of various fields. There is a name. The name doesn't have to be the name of a natural person; it can be a title or a role. There is the organization name, there is an address. Here, it can be postal and often is. An email address, a voice number, and a fax number. I'm not sure how much longer the fax number is going to survive. I suspect that's a relic of another era, but for now, it's still a record we have on the books. Next slide.

Here we have the name servers listed on the website. This is sort of one of those technical facts. It's a direct reflection of the NS records and the glue records that appear in the domain name system itself. Strangely enough, not through anyone's design, we don't actually publish the DS records on the website. We do via the other mechanisms that we use. For various uninteresting reasons, we never put the DS records here, and only one or two people have ever asked, so that's the current state of affairs there. Next slide.

Then here's what I guess I'd summarize as registry facts. These are user-supplied. Each ccTLD manager can essentially nominate what they want to appear here. We have the URL for registration services. Here, it's at the discretion of the registry manager what makes

sense. All we do is validate that the website is contactable, like we get an HTTP response if we connect to the URL, but we're not validating what is actually on that page. There's no requirement as to what kind of information a ccTLD manager puts there. The WHOIS server, again, we're looking for technical accuracy that there is a port 43 WHOIS server at the other end, but we don't validate or have any requirements as to what kind of information that WHOIS server supplies. The WHOIS server can literally say nothing, and that would be deemed acceptable. For the RDAP server, there is a minimum level of compliance testing here that we send a query and we expect a well-formed RDAP response back. But we, again, have no requirements of what kind of information is actually in that response. Next slide.

Then here is the IANA report section. This is where, if there is a report on delegating or transferring a TLD, we would just link to it here as a courtesy. We're thinking about adding the LGR tables as well, so if that particular TLD has registered label generation rule sets, this would probably go right about here. Notably, it doesn't include any other resources: no third-party resources, no ICANN resources. For example, it's very common that ICANN has an accountability framework or an exchange of letters with an individual ccTLD. Since IANA is not a party to that, and this is just strictly IANA-related references, we don't link to that on the ICANN website or anything like that. Next slide, please.

At the very bottom, I can't see it; it seems this page is a bit zoomed in. Imagine just below the bottom of the screen, there are some

dates. There is a registration date and a last updated date on each record. This is a source of some confusion because these dates are actually database record dates that don't necessarily tie to any particular event in real life. Particularly the registration date; it's the date that it was put into the IANA database, which many people assume is the date that TLD first existed, which is not the case. Retroactively, we have changed this date in some cases to try and better match with the date that we understand it being created. We're actually looking to evolve past this a little bit. We're undertaking some work right now to add a new field which will literally be the registration date for that purpose. The idea would be to have an explicit delegation date and somehow signal our confidence level. If we know that is actually the date and we have it in our records, we will signal that this is actually the official date it was put in the root zone for the first time, and you can take it to the bank that this is the day that TLD was created.

For some older TLDs, we simply don't have definitive records. For some we do, but we don't have a complete record of all the decisions that were made prior to ICANN's existence. We're thinking we'll have some kind of confidence level or something that will be tied to that, where we can say, for example, it's an estimated date or a time range. We're still exploring exactly how to signal that, but we're looking at doing something along those lines. When we do that, I think we might deprecate this particular field just because it causes so much confusion that it's there. Next slide, please.

Some of the context here is we're actually redesigning this page. The redesign is sort of tangential to the purpose of the fields. The actual fields and everything I just told you are the same. We're not changing our record-keeping processes as a consequence of this. This is just more of a user design exercise to make it a bit easier to use. Here, we are planning to add a few extra fields we don't have today. We do have them in our internal database; we just haven't exposed them yet. The most obvious one is we have the language and the script for IDNs, and it's not something we currently share. We're going to add that to the record-keeping. Another thing is that we have some little information icons there. You can see it against a label at the top. This is just some supporting definitions for some of the fields where it's lacking, to help people understand what those fields are actually for. Next slide, please. All right, here we go.

Here are just some additional comments that came to mind to stimulate discussion. I think one thing that we still see as a misconception today is something that we did in 2022. For those that are very close to the ccNSO, you will know this. Going back many years, those admin and tech contacts in the database served double duty. One was to be a public point of contact, but they were also required to cross-validate every change for that TLD in terms of our operational procedures. The admin and tech contact played a very significant role operationally. But in 2022, we disconnected those two responsibilities. Today, every ccTLD can manage their own set of private contacts with us who are authorized to do what on their behalf. It doesn't need to be two; they can specify as few or

as many as they want. They can specify granular access permissions. The point is that the public records now have no connection to those responsibilities. The public records now are just that: public points of contact that have been nominated by the ccTLD manager, and it's really up to them what they specify there. What that's done, for example, is help ccTLDs who want to nominate a help desk-style point of contact to be in the public records, whereas privately, they want someone in a senior position of authority within the organization to sign off on changes to their ccTLD. It's really separated those two areas of concern into different kinds of contacts. To be clear, those private points of contacts are private, so they don't appear in any of these public records that you're seeing.

Another thing I would comment is that generally our objective is to maintain consistent operational procedures across all types of TLDs. The root zone is a commons for multiple TLD types. Our systems, like our root zone management system, are used by gTLDs and ccTLDs alike. For that reason, we're trying to have some consistency in approach so that we don't have completely different interfaces and fields depending on TLD type. As we get further into this and we look at potential changes to our approach, I know that I'll be sitting there with a mind to the potential impact on other TLDs. That's not to say we can't customize, and if there is an imperative to have some specific approach that's just for ccTLDs, I'm sure we can accommodate that. But I would like to make sure that the work is worth the outcome. We're not just customizing it

for no particular reason, but for a compelling reason that outweighs the desire for consistency across all TLD types.

I mentioned before, IANA does not publish the country's name in connection with an ISO Alpha-2 code. That said, when you enter into an address field, we do have a dropdown of all the ISO country names. You can pre-select one of those, but you can also override it. If you want to call your country something else, you can go in and override the default setting and select what you want. There's a competing tension there. On one hand, we like to see consistency and have a common set of country names, but by allowing people to override the country name, that's introducing some customization that divorces us from consistency. I will note that some of the criteria we have for how we assess or limit these fields today is current practice. Some of our data is 30 years old, so the standards that were applied 30 years ago are obviously very different from today. This speaks to some of the core issues that this group will be looking at, which is compliance or consistency. Even if we have a good set of operational rules today, how do we bring all those historical records up to speed with modern-day requirements? We have some bad data in there, or maybe it was good data at the time, but it's not consistent with current practice. Hopefully, that's a useful tour through what we publish for ccTLDs.

JORDAN CARTER

Thanks very much, Kim, for that presentation. We've got around about ten minutes for questions and any comments. I did have one

question in the chat from Stephen Deerhake. He was asking, 'Will the JSON data attributes include private contacts that ccTLDs have submitted to ICANN via the RZM?'

KIM DAVIES

That is the easiest question I've had all day. Absolutely not. The private contacts are private; they won't be in public data.

JORDAN CARTER

The private contact structure is out of scope of this working group, which is about the public records. Thanks for the question, Stephen. I've got Peter Koch, who has a question.

PETER KOCH

Yes. Thank you. This is Peter Koch, DENIC, and a member of the committee. I have two clarifying questions on that final slide. The first one is on the first bullet item, or the sub-item actually, this welcome opportunity to have a specified address there to follow public demand. First question is asking for confirmation: is that true for both the admin and the tech contact? Okay. Do you apply any rules about these contacts going back to the requirement of being in country? If I delegate just for the sake of the argument, if I want to outsource my help desk to somewhere else, would that have to be in country, or is that rule applied on a totally different basis?

KIM DAVIES

It is applied. This is one of those areas where I think the framework of interpretation favored this, which is there is a literal interpretation that is hard to avoid: RFC 1591 explicitly says the administrative contact must be in country, and this is the administrative contact. Absent superseding guidance to interpret it in a different way, that's not something we waver on. Now, part of the discussion is whether that should evolve. But as of today, we assess that even if it's a nominal help desk point of contact, it still must be in the country.

PETER KOCH

Thank you. But it doesn't have to be the same entity. It could be some other entity, and you would just look at the country part.

KIM DAVIES

Correct. It's a point of contact that's been designated by the ccTLD manager. It doesn't have to implicitly be personnel of the ccTLD manager; it could be a third-party entity.

PETER KOCH

Okay. So for further consideration, and not now, we might want to look into what the real purpose is here, because I understood that once these private contacts exist, there is no use for IANA of this information. Is that a correct interpretation, or am I going astray?

KIM DAVIES

I think that's fair. In terms of IANA's operational procedures, we no longer consult that information. But that's not to say we don't have procedures around making sure it's kept up to date and accurate. We test the email addresses, for example. You can't just throw a throwaway email address in it. But it has fundamentally shifted to be a resource for public consumption, as opposed to a resource for IANA to operationally use.

PETER KOCH

Yeah, thanks a lot. I have one other clarifying question, but I will let somebody else go first.

JORDAN CARTER

Thanks, Peter. That's very pro-social of you. I have Nigel next, and then I have Irina.

NIGEL ROBERTS

Thanks. I don't want us to get into the weeds on this. First of all, Kim's absolutely correct. Like many of us who were in this room, we spent over a decade and a half on various FOI-related working groups. He's absolutely correct. The literal interpretation of what's in RFC 1591 is that the admin contact must be in country. There are some other spins that you can look at. There was a very notable ccTLD which, if I remember correctly, was delegated to the University of Puerto Rico, but it was actually .kn twenty years ago. The point is that this is what this group's all about, discussing these things. The ccTLD manager could quite conceivably be

incorporated in a different jurisdiction, just as the Trump Organization in New York is probably incorporated in Delaware and possibly other places as well.

JORDAN CARTER

Thank you. That's the first time I've heard that name on the record today. Irina.

IRINA DANELIA

Hello, everyone. Irina Danelia for the record. I'm with .ru, and I'm a member of this study group. Just for the broader context, I would be interested to know what are the non-public records in the IANA database. Secondly, it was mentioned that the policy or the procedure allows different types of contacts to make different changes in the IANA database. I remember several presentations on this topic, but I assume there is a resource where this is described in detail, and I think it would be also interesting to look at.

KIM DAVIES

In terms of our private records, there's a lot that's not public. In terms of the Root Zone Management System itself, it's predominantly just what we've already mentioned: we have a regime of private contacts. Above and beyond that, we have all the transactional data, like what requests you've submitted, status of requests, and all those kinds of exchanges. There's credential information and all sorts of stuff to support that. Beyond that, we

have a ticketing system where all of the exchanges that relate to the ccTLD have been recorded. We've been using a ticketing system for root zone management since 2005. Every email transaction pertaining to your ccTLD we have in a ticketing system. Obviously, we have the working records of when we're doing assessments on things like delegations and transfers of ccTLDs. There's a lot of records. It's not structured in the sense of sitting in a database, if that's the question, but there's a lot more that we don't publish than what we do. Does that answer that part of the question?

IRINA DANELIA

In our registry database, for example, we also have a lot of data, and some of it is public, and some is not public. In general, we can split this data into several groups, like data to identify the domain name holder, data to contact him, and technical data associate which is needed to make this domain name in our case working. There is also the ticketing system and logging of the acts, et cetera. But at least there are main types of information, and some of it is publicly available while some is not publicly available.

KIM DAVIES

I think it's much the same. Without going too much down the rabbit hole, we have a bunch of additional data, certainly in terms of those authorizer contacts. We have all points of contact there. We have emergency points of escalation. There's the TLD Ops Security Initiative, and those that participate get that bi-monthly list of points of contact; we save that within our team so that we might

need an emergency response. There's all sorts of information we have that we don't necessarily publish. We have some monitoring tools where we're monitoring the apex of each TLD, and if we see a disconnect there that is operationally urgent, that triggers us to do some out-of-band follow-up. This is not something that's in the policies, but it's just good operational hygiene stuff that we do. We have a bunch of different stuff there. If you wanted a more comprehensive inventory, we could probably explore that, but it's not something we'd ever considered making public as part of the core IANA obligations, which I think is where my head was initially when looking at this problem space.

JORDAN CARTER

Yeah, I think the focus here is on the public records. But because we don't know what's in the private records, we can't make any judgment about whether there's anything relevant to the public records within them. To the extent that there might be, or that you could provide some kind of overview material, that would be helpful, but let's not boil the ocean with it. It's fortunate that there are no more hands up because there is no more time. Oh, there is a hand up. Peter, is it a quick question?

PETER KOCH

We can move forward.

JORDAN CARTER

Okay. Feel free, members of the study group, to use the email list to ask any follow-up questions. If we could go back to our slide deck, Claudia, we will finish on time. The only other slide really to talk with you about is the upcoming meetings of the group, which is of most interest to the people who are members of the study group. If we go to the agenda. The next page. Up, next page. And the next page. These are the planned meetings. We're going to try a fortnightly cadence, and we're not doing a Doodle poll to find the time. This is the time that the group will meet in that five-meeting phase. If it's horrible for you, I apologize. If it really is horrible for you as a member of the study group, please ping me, and we might need to go into a rotation. That's the first pitch. If you're a member of the group and you have any problems with this timing, we'll discuss it at that first meeting. Inevitably, that might be horrible for you as well, but if it is, let me know. You will survive. And we did say there was a moment for any other business. So I'll just ask briefly, with less than three minutes to go, if there is any other business. Assuming that there isn't, I'd just like to thank you, Kim, for producing at relatively short notice this exploration. I hope that it started to give people a sense of the phrase IANA Public Records. Now you're starting to come to grips with what exactly those are. Hopefully that was a useful start to the work. March 25, what's the more granular plan to advance that understanding and research on our part? Thank you all very much for coming along for this session. I hope you learned something and/or contributed something. For those of you on the study group, we look forward to seeing you at the next meeting. There will be a brief update on what happened

here today in the ccNSO members meeting on Wednesday, but there are no other meetings of this group during this ICANN 85. So with that, I hope you've had a lovely day minus one of ICANN, or I don't know if it's day minus one or day one or day zero, but enjoy the rest of your Saturday and see you around the traps. That's the end of the meeting.

[END OF TRANSCRIPTION]