ICANN85 Mumbai | CF – Customer Standing Committee Work Session
Monday, March 9, 2026 – 9:00 to 10:00 IST

CLAUDIA RUIZ

Hello and welcome to the customer standing committee work session. My name is Claudia Ruiz and I am the participation manager 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 put them in the chat for your reference. During this session, questions or comments submitted in the chat will be read aloud if put in the proper form as noted. If you would like to speak during this session, please raise your hand in Zoom. When called upon, the virtual participants will unmute their microphone.

On-site participants will use a physical microphone to speak and should leave their Zoom microphone 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 Rick Wilhelm, chair of the CSC. Thank you.

RICK WILHELM

Very good. Thank you, Claudia. Good morning, good afternoon, good evening to anybody online. If you're in the physical room,

please join the Zoom. Normally, we're laid out in a U. We couldn't find a bigger room. We tried. We tried. We tried.

CLAUDIA RUIZ Rick, can you put the microphone a little bit closer to you?

RICK WILHELM Oh, get a little bit closer. Thank you. Yeah, I hesitate to move them when they're on the dais because otherwise the AV team will get mad at you. Yeah, we couldn't find a bigger room. Usually, we're laid out in a U, which, as I noted before, I much prefer. But c'est la vie. Okay.

So, welcome to Meeting 101. When we get going, I don't think we're quarried. We are quarried. Okay, we are quarried. Okay, very good. Oh, excellent, we're quarried now. Thank you, Bart. So, why don't we go ahead and get kicked off because we've got work to do. We've got an hour and 15 minutes today. So, why don't we, Bart, you want to roll us through our action items?

BART BOSWINKEL Yeah, it's pretty straightforward. In principle, the action items are completed. Some are referenced on the agenda item. So, back to you.

RICK WILHELM Okay, very good. If anybody has any questions about the Action Items? Nope. Okay. So, now we're going to go to the January, let's

see, this says, I guess it's the February report, a minor typo that I failed to pass the test. Hiro let a bug through to see if I was awake. And so, we do this to each other. So, when the agenda was circulated. So, I'm not sure who's going through the, okay. So, that's going through. Okay. Very good.

SELINA HARRINGTON: Thank you. So, in February, we met 98.4% of the naming SLAs. We missed one SLA, which is the technical tech technical check retest for routine technical change request type. And the summary of the issue was the test duration was affected by unreachable servers encountered during validation of consistency between glue records and authoritative data.

While resolving one address for a TLDs name servers, a referral was returned that included multiple glue addresses, all of which were unreachable due to the two sequential queries and the preference for glue records over external name servers. Several checks exceeded the target time of 10 minutes.

RICK WILHELM: Very good. I believe this is a failure that we've seen before. Is that correct? Or is this a novel failure?

SELINA HARRINGTON: I believe it's something similar to what we've seen before where it's on the retest that we end up failing the SLA. The initial tech check timing was almost the same as the retest, but because the SLA is

different there, we didn't, we didn't fail the first tech check, but yeah, because of unreachable name servers, we've had this run into this issue before.

RICK WILHELM: And did the customer fix anything or did they change anything in their DNS or did the network reachability change? In other words, did they, when it finally resolved.

SELINA HARRINGTON: Let me just double check. So it kept referring between two unreachable name servers until another one returned a referral to a reachable host, which allowed the test to complete ultimately. But there were several retests that I think, I believe timed out initially.

RICK WILHELM: Okay. We've got two hands in my queue. I think I've got first Gordon and then Peter. So, Gordon first, please.

GORDON DICK: I wanted to ask around the -- obviously reachability is something that is related to both networks, your network and somebody else's network. Where are the tests done from one single point or are they done from multiple points in that process? I'm just wondering how much of that is a failure, is the reachability a failure of the TLDs name servers versus interactions of the internet with your tests?

SELINA HARRINGTON I believe it's done from multiple points. Kim, is that correct for the technical checks?

KIM DAVIES Yeah, we have two sites we do them from. That said, it's been on our backlog to add a more diverse testing location. That said, I think if it's failing from two sites, US West Coast, US East Coast, I think it's usually evidence there is some kind of problem there, even if it does resolve in other parts of the world. So it usually merits a dialogue with the TLD manager regardless, just dig into what we're seeing and see if it's more on our end or their end. But at least we know there's some issue there.

GORDON DICK Yeah. No suggestion that there's no need for dialogue, it's just to understand as people look at it. It's certainly something we've seen evidence of failures between those tests coming from your network to not our DNS network, but a third-party DNS network, which has caused confusion because even from multiple sites around the world, we've had challenges to see any problems.

RICK WILHELM Thank you, Gordon. Peter, your hand went down. Did your question rhyme with Gordon's?

PETER KOCH

Yeah, thanks. This is Peter Koch. Thank you for giving the mic to an observer. Yeah, that was the direction, but the follow-up question, because we had that in the IFR2 review a couple of times that these failures might be booked to the wrong party. And the question is what to do about that. And I understand that maybe investigating the issue further might be more effort than just accepting the bad, the missing the SLA. But for the next IFR review team, they'll have to go through the same thing and then doing that again. So, it might be good to fix it.

RICK WILHELM

Anil has a hand up. Anil, please go ahead.

ANIL JAIN

Thank you. Anil, for the record. I personally feel that this problem of non-reachability of the server is not a bigger issue. And this, with the situation of the networking, which we are able to see today, it may come again. So, I personally feel that it is a minor issue. We should not take it very seriously. Thank you.

RICK WILHELM

Thank you, Anil. And I agree with Anil's point. And I think that Gordon was also not saying that this issue is that major. It's really the question of the -- and I'm not sure about exactly the quite correct word, the architecture of the test network. And I like how many observability points do we have for testing these things?

And Kim said, well, one east and one west, which gives us two. And I think that perhaps prior to this, Gordon perhaps thought that, well, maybe there's only one. And then if you've got only one point of observability, then that would indicate a head scratcher, technical term that we can all probably deconstruct.

And so, but two, U.S. East and U.S. West, as Kim said, well, two is not three, even for sufficiently large values of two, as one of my math's professors might've said. But if you've got problems from two, that I think the transcript would probably -- I think Kim said, that shows a problem somewhere. And assuming that these two observability points aren't single-tracked, somewhere ridiculous, then it probably indicates there's a problem somewhere. So, might deserve a little bit of further study, but I think it probably knocks down one question that Gordon had first. So, I think we're kind of generally okay. Fair enough. Did I deduce, boil that down? Pablo, please go ahead.

PABLO RODRIGUEZ Thank you, Rick, and greetings, all. Tacking on that idea, but Gordon, correct me if I misinterpreted what you meant, but I believe I heard that while what Kim told us, testing from two sides is, there is something that's happening, right? So, many flies can be wrong. Something is happening someplace. But I believe I interpreted that you said that when it happened to you, you were not able to find a problem with your systems. Was that correct?

GORDON DICK

So, it was name servers that were not our own network. We have monitoring on them though, and we were not able to identify an issue from the global network. That being said, I still think as long as there's no single point of failure in your reachability, I think that is a reasonable scenario.

Part of my concern around it was actually around the messaging that goes out to all the contacts, because there are different levels of understanding of the people that approve it, and it causes levels of concern that are not necessarily relevant for the failure, because TLDs were operating absolutely fine. We actually got some details from one of your team, but it was still difficult to determine what the issue was. But I think it's reasonable if you've got multiple points of tests.

RICK WILHELM

Yeah, that's an interesting point about how, and I don't know what that messaging is, right? But the problem that Gordon is pointing out is a slightly different one. You know, it's sort of a difference between, like you didn't answer your phone, and when a phone rings to voicemail, you really don't know is the phone not on, or is the person not picking up, or is the person asleep, or are they just not ignoring you, or do they have do not disturb, or are they on another call, or just all you know is, I did not connect with that person.

Poor analogy, right? But it's an interesting challenge. And as Gordon says, when those alerts go out, the technical person might be like, look, the name server is running fine from our end, but it doesn't mean it's running okay from everywhere on the internet. So yeah, interesting point. Thank you.

And interestingly, and oddly tied into our discussion about SLAs. So, that's why I sort of, on a normal meeting, we might like, we might kind of shut this discussion down quicker, and be like, it's okay, no big deal. But today, I kind of purposely let this broth simmer on the stove a little bit. So, we could percolate our brains here. Thank you. Stale hand, Pablo? Stale hand. Okay, very good. Thank you. Don't turn myself off. Okay, very good. Thank you. Looks good. What's next, Bart?

BART BOSWINKEL: Next one is the findings, your findings. Oh, okay. Findings. Jennifer, please.

JENNIFER BRYCE: Yes, that's me. And I have something in my eye, which is called, I'm not that upset about the SLA being missed.

RICK WILHELM: It's a very emotional topic. I mean, it's always very touching when I see these.

JENNIFER BRYCE People are concerned for me. But yeah, just the standard report here that you can see on your screens that reflects the 98.4%. And then the standard language at the bottom there, the missed SLA has been satisfactorily explained by PTI. And the CSC has determined that this exception is no cause for concern. No persistent problems were identified and no further action is needed. So, I'll turn it back to you, Rick, if you want to call for approval. Thank you.

RICK WILHELM Very good. Thank you very much. I approve. Bart, you can -- well, we'll just kind of go around the room. Thank you. Rick approves.

BART BOSWINKEL The question is easier. If you are opposed or no opinion, please raise your hand. Otherwise, we'll approve. It's approved.

RICK WILHELM And any abstentions or disapprovals? All right. Approved by unanimous acclamation. Very good. All right. Good show. What do we have next?

BART BOSWINKEL I don't have the next one is what we briefly discussed on the last meeting was the comment on the bylaw changes and charter changes, section 17 and 18.

RICK WILHELM Very good. We've had this comment. I think this has been out on the list. It'll probably pop up on the screen here.

BART BOSWINKEL Can you open the document please?

RICK WILHELM In a second? It's a very short and sweet comment that generally says we're all good with the proposed changes. It has a clarification that you can see in the bottom of paragraph one that says for avoidance of misunderstanding this we've got a little bit of a typo. This is the ccNSO one. I think we've got a wrong tab. Our CSC one is similar. It's probably a search and replace different. That's the advice.

BART BOSWINKEL Can I share the document? It's in the link.

RICK WILHELM We'll get there in a second. So our comment looked just like that one, only instead of ccNSO, it said CSC. Okay, I'll be able to share. Oddly similar.

BART BOSWINKEL Just a minute.

RICK WILHELM

We'll get in a second. No problem. It had a bit in it about that this is a position of the CSC as a group. And here we go. Good show. Yeah, for avoidance of misunderstanding this comment is a CSC comment only. You can see where the cursor is blinking there towards the bottom of paragraph one. And the views expressed in this comment should not be interpreted as abuse from any of the members individually nor from any of your member liaison appointing organization.

That means that if since I came here via the RySG, the RySG can say whatever it wants about this and it doesn't really bind me to that. So, we fully support that stuff. And then we've got a bit at the bottom to avoid some scenarios. We guide them to avoid some overlapping things because that would put stress on the staffing of the various reviews.

Peter nods because he, and to a certain extent I have been in a situation where we've kind of gotten into these reviews and it's better to skip from one to another than be doing both at once because that places follicular stress into our lives. And neither of us are the age when we need that. So, that's sort of it.

The rest of it is pretty much boilerplate. And then we've got a bit in the middle there where we make it clear that we support and we thank them for the work and blah, blah, blah. Anybody have any questions? This was out on email and so I'm just filibustering here

for a minute to allow everybody to reread it and reload it into your RAM.

Questions, comments, concerns? Anybody have any problem with us publishing this thing? Thank you, Anil, for the thumbs up. All right. We'll move on from that. So, Bart, you can go ahead and publish that. This is ahead of the deadline by a week or so, maybe two weeks. Look at us, setting them up and knocking them down as they say. Okay, good show. I think our next item is vice chair elections.

BART BOSWINKEL And so, the first step is maybe, Hiro, you can lead this bit. The question is, first of all, are you available to stand again, Rick? So if you are, you're not term limited.

RICK WILHELM Yeah, I'm not term limited and my employer has not told me that I should not or may not run. So, I'm okay.

BART BOSWINKEL So the question is, is there a nomination for a chair? Who wants to nominate a chair? Should I nominate Rick then? Yeah, Rick. And do you accept, Rick?

RICK WILHELM Yes, I accept. Thank you. Appreciate it.

BART BOSWINKEL I have a vote from the members. Do you support Rick? And I assume you're abstaining, Rick. But the other three members, do you support? Yeah, yeah. Pablo. Okay. Congratulations, Rick.

RICK WILHELM But I didn't have to spend any of my campaign money. This is so fantastic. Wow.

BART BOSWINKEL Maybe the same, Hiro, you're not term limited. Are you in principle available to stand again?

HIRO HOTTA Yes.

BART BOSWINKEL Is there a nomination? Pablo, nominates. Do you accept the nomination, Hiro?

HIRO HOTTA Yes.

BART BOSWINKEL May I have a vote of the members?

GORDON DICK Aye.

BART BOSWINKEL Congratulations, Hiro. That closes item number 4B.

RICK WILHELM Very good. I'm glad that vice chair election went okay. Otherwise, I was about to resign. So, good show. Thank you, Hiro. Glad to run it back. Okay, good stuff. Is there anything that we have to do or we have to probably publish something, Bart, is that correct, that we're re-upping?

BART BOSWINKEL It will be in the cover note and that's it.

RICK WILHELM Okay, good show. Very good. Thank you, everybody, for that. And we're glad to do this again. Good stuff. Are we now diving into the SLAs? All right, very good. The main event, the undercard is over. No more faffing about. So now we're going to get into the SLA stuff. We talked about the review grid. We talked about the SLAs. And I don't want to, we'll take that as read, as we would say in some Board meetings and stuff that I do.

And so, why don't we just kick on over and dive in and make sure that we get ample time to sink our teeth into Kim's presentation, which we did not do at some meeting that I can't -- which meeting was it, Kim? The last one? Or was it…? Oh, Dublin. So, it was the

last one that we were in person. It seems like longer than that ago. I don't know. Man, my memory.

BART BOSWINKEL

Okay. Just for the record, we had about a year ago when we started this. And then the intention was to really start the discussion and kick it off in Dublin. But there was a lack of time in Dublin because we ended up with, I don't know, I can't recall the other topic. I probably talked to you. You just had five or 10 minutes.

RICK WILHELM

Yeah. So, we're going to dive in and have Kim. Yeah, that's right. There's the Dublin building. So, Kim, further proving that Kim was right. So, he didn't even change these slides. It's like there it is, the ICANN84. Okay, very good. Thank you, Kim.

KIM DAVIES

Yes. In preparing this meeting, I thought I was crazy because I thought I'd done slides before. And sure enough, I hadn't. I just never presented them.

So, this is a high-level summary of IANA's thoughts on this. Actually, I'm going to move closer because now I think what I did do in the Dublin meeting is the first couple of slides that just kind of introduced the history about how the SLAs came to be. I'm happy to present them again, but in the interest of time, do you want me

to skip ahead a few slides? Get a feel for the chair's prerogative. Hearing nothing.

So, let me just quickly go through these slides then. So, the SLA format and structure that we have today is largely a product of the IANA Stewardship Transition. One of the design teams was focused on creating the SLAs. There was a sub-team of three people. I was one of those people that was involved in authoring the original slate of SLAs.

They've since been adjusted a little bit. IDN tables or label generation rule sets was a notable admission that was recognized in 2019. So, they were added then. And then also it was recognized there were some challenges with the way ccTLD delegations and transfers were being measured. So, there was modifications in 2020 to make them a better measure of performance there.

Another structural change that happened was originally the SLAs were embedded in the contract, which meant that to change the SLAs required a contractual amendment, which is a difficult process. So, they were converted into a standalone document that is cited by the contract. And this allows us to update the SLAs more freely in coordination between PTI, ICANN and the CSC.

Next slide. So, what is measured for root zone management? Essentially, we divide all the measures into five categories. There are changes to the root zone that do impact the root zone file. These are technical changes, NS records, DS records and the like.

There are changes that do not impact the root zone file. So, this is whenever those technical elements are not changed.

So, often this is changing points of contact, things like that. The notable difference between these two scenarios is that if it's not changing the root zone file itself, we don't have the technical check procedure that we just talked about earlier in this meeting. And we also don't involve implementation in the root zone, which involves working with VeriSign as the root zone maintainer, promulgating the changes to root server operators and so forth. So, a lot of the measures of that tail end of the process don't make sense if the root zone file is not changed. So that's why we have a distinction there.

Delegation and transfer of a gTLD is the third category here. Delegation and transfer of a ccTLD is the fourth category. Why do we measure Gs and CCs differently? Because they're fundamentally different work processes for IANA.

gTLD delegation transfer is much more straightforward, whereas for ccTLDs we're for doing significant due diligence and communication with the customer, usually over a period of many months. So, it's just a qualitatively different process.

And then lastly is sort of a catch-all for everything else we do that is not routine. An example of a very non-routine change request that will happen from time to time is updating a root server. So, an IP address, for example, of a root server might change. So because they are so novel and so rare, we just have them in their own

bucket, which there's no actual SLAs attached, but there are reporting on processing timelines.

And then separate from root zone management we have label generation rule sets, also known as IDN tables. These are measures around how quickly we process changes there. We have SLAs on the SLAs and how quickly we process them. We also measure the availability of SLAs, if you will, by making sure that we deliver our monthly reports on time.

We have availability of our SLA dashboard measured to make sure there is a real-time component to this, where at any given moment you can go in and look at our SLA performance and interrogate that. We also measure inquiry response. So, if you ask IANA questions about the root zone, we measure our time to respond to those in terms of just general inquiry processing. And then credential changes. So when you make a change to your credentials in the root zone management system, we measure how long it takes for those changes to be made.

Next slide, please. So, what kinds of measurements do we make? Predominantly they're time-based. So, we perform a certain activity from beginning to end within a certain period of time, measured on the clock. And then there is an SLA for percentage compliance. For example, do this particular process 95% of the time within X minutes. And so, that's the way most of the SLAs are structured.

We also have measurements for accuracy, what percentage of requests were actually implemented correctly. For system availability metrics, it's the percentage of time the systems are operational. And as I mentioned, some of the metrics are not associated with a commitment. So, we do report on them, but we don't say we will do them within a certain amount of time. They're their advisory to the community.

So, the community can monitor trends and make observations and, still take corrective action or notice a trend that is worth investigating. But it's not something on a month-to-month basis. We're discussing with the CSC whether either passed or failed. Next slide, please.

I think you're all familiar with this view, but this is aggregating sort of the measurement category with those five categories I mentioned earlier. It sort of shows all the elements of what I just talked about.

All right. So, that was the summary of what we do and how we got to where we are. This is sort of the meat of the request I had from Rick and Hiro, which is like what are IANA's thoughts on the current SLAs. So, I think this is perhaps the substance of what I wanted to share.

So the first thing, and this is in discussion with my team, again it was just in the sort of weeks prior to Dublin that we explored the issue and had some preliminary discussion and put this together. But the first thing we observed is that some SLA categories don't

actually measure anything. While some of them measured something back in 2016, our operational practices have evolved. And so, it sorts of obviated some of them.

And frankly, some of them never really measured anything even back in 2016. But there was a philosophy within the design team that was building the SLAs out that there not be any aspect of the timeline from beginning to end that was under IANA's control that was not measured, no matter how material that phase was. And I think it was a notion that if it wasn't measured and IANA didn't perform there, then there could be ultimately a lapse in end-to-end processing time that wasn't discussed or accepted. But yeah, so that was the philosophy that led into that sort of structural design of the SLAs.

So as of today, the way we operate, acceptance recognition this is essentially measuring from a world where people would submit a ticket and how long it took for them to get the ticket acknowledged, the email back saying we've received your request, here is your ticket number.

And so, that was the way it worked back then, both for root zone management submissions in our system, ones that you would interactively submit, you would still get an email back as your confirmation that it was submitted in the early versions of that software. But more importantly, there was no barrier to you just firing off an email to IANA, like there was no obligation to use the system.

So, many TLDs would just write an email, dear IANA, here, I just want to add this NS record. And they would email it to us. And so, we would email back and say, okay, this is your request, here's your ticket number, and we're starting to process this. So, it was measuring that latency there.

Now today, it is a requirement to use the root zone management system. So, you must log in interactively to submit a change request. And as part of that workflow, when you submit in real time on the web interface, it says your request is submitted, here's your ticket number. Yes, there is a follow up email as well, just for record keeping purposes, but it is a real time transaction.

So essentially, it's always zero, because measuring it is the time it takes to sort of, technically do a database insert into our system. It's just, it is a millisecond level transaction that's happening here. It's not potentially minutes, slash hours, slash possibly days, as it was when this was first conceived.

Similar observation on the recognition of confirmation. This is, again, a situation where there were manual workflows, where a TLD manager who's asked to confirm, do they want to make a change as part of our workflow, part of that verification that it's authorized by the correct people at the TLD manager, that we would ask them to confirm, do you consent to this change?

Once they'd supplied their consent via whatever form, how long did it take for our team to read that consent and record that consent? So that's essentially what that was measuring. Again,

today, with automation and a constraint that we've applied in recent years, you must log into a system and authenticate. And when you supply your confirmation, it is through the interface and it's recorded in real time.

So it's, again, milliseconds to insert into a database and you can see in the reports we've supplied in recent years, it's effectively measuring nothing of significance.

Time for requests to be lodged via email. I already touched on this a little bit, like 2016, as long as we could read it, we were accepting of whatever, if you wanted to send it via phone, fax, email, postal, if you knocked on our office door and said I want to make a route change, we would probably entertain it. But now we do obligate customers to authenticate via the web interface and submit, except in rare cases, like the rare cases are essentially, if you're doing a ccTLD transfer, you're not a customer of record today, obviously we accept that. That, by the way, would fall into category four, which doesn't have an SLA on it anyway. Or, emergency disaster recovery scenarios.

We talked a bit about that yesterday in the ccNSO, where, the power's out, like complete infrastructure devastation and you tell us you can't even log in because we have no connectivity and we're trying to restore things, of course IANA's going to support that kind of transaction, but it is so rare, I would suggest it's not relevant for the purposes of routine SLA measurement.

Time for authorization contacts to be asked to approve change requests after completing previous processing phase. That's a mouthful, but essentially, again, that was the philosophy that if there were gaps in the end-to-end process, maybe IANA would just sit on its hands for a while and it wasn't measured by some way, so therefore it would not get captured. But today it's all automated, the next once the conditions of one phase of processing are satisfied, it automatically moves on to the next phase. There's no human intervention or anything happening there.

And then time for response to be affirmed by contractor. This is, probably feels like legalese, but essentially, it's like how long does it take for us to confirm that your resume change is being implemented. Again, it's all fully automated today. There's really no delay or role for human intervention in this process.

So these are examples of like, I think, well-intentioned SLAs that, we can continue measuring, but I think in the context of looking at a holistic review of SLAs, you're never really going to get any deviation in the numbers here. And I think the question for us to explore is, is there value in us reporting these? Are these useful measures? And I'm going to get to this later, and probably a caveat for everything we talk about here.

For me, there's not a lot of additional effort to leave things the way they are. Like, I don't have to change my software, our reporting software, it's all automatic for us. But I think there is a competing concern like, is the CSC focusing on the right measures? And if we

look at the suite of 63 things we measure today, are these the priority? Are these of equal priority to the others? And perhaps the answer is, no.

Next slide, please. Technical check categories aren't calibrated to the time and recurrence. So let me explore what I mean a bit about this. So currently, in the lifespan of a change request, we will perform the suite of technical checks, and these are the ones for name servers, multiple times throughout the course of a transaction, at least twice.

The first time we check, like right after submission, we call the first technical check. And the design assumption in 2016, when the design team was creating the SLAs, is that in the normal course of business, when someone submits a change request, they're most likely to have errors at the beginning, because it's their first time submitting it.

And if there might be something they didn't think about, our system will catch it on the very first pass, report to them there's an issue with the technical configuration, and in normal circumstances they'll go, oh there's a problem, go remedy it, click retest, or there's an automated retest I think every three hours, and therefore and then it will get resolved.

So, the assumption was that the SLA for that first technical check should be less aggressive than the others, because it is more likely to catch configuration issues, and of the kind we talked about earlier in the call, like timeouts, things like this, that might make

that test run take longer, and therefore it was appropriate that IANA not be held to the highest standard for that first pass, and that the SLA is more generous at that time.

So that was the thinking. But the reality is the opposite, because if you think about it, and it makes sense in retrospect, yes, everything I said is true about the first tech check, but if you pass the first tech check, we do retests, and so that's the second category. So, who gets retested? Well, only people that failed the first check.

So, the first tech check represents the collection of people that passed it on the first instance, of which time the checks usually just go very fast, and those that have a misconfiguration, at which point the tech checks typically take longer. But then the retest population is only people that failed the check before, so they're always going to take longer.

So the population of retests is always, on average, going to be a longer set of tests for that reason. So in a sense, the intuitive idea that the first tech check would be slower than the retest, it's actually the inverse in reality, that the retests take longer than the first tech check. So, that was something we learned.

Now, in talking with our team, perhaps we're not at the solution phase yet, but one of the suggestions we threw out there for consideration is, before I get to that, let me just explain what supplemental is.

So, the supplemental tech check is the second round of technical checks we do, and we do that, so we do the first technical check right at the beginning of processing, and we do supplemental tech check right at the end, just before we pass a request to VeriSign. And why do we do this? Well, it is a bit of a legacy from when root zone processing often took days or weeks, and there could be a deviation, like what could have been the TLDs configuration on day 1, by day 28 or whatever it is, when we're ready to put it in the root zone, maybe something broke in the meantime. And that wasn't entirely academic, it was happening more often than I think anyone would like to admit.

So, that's why we had a process, like just before we dispatch it off to the root zone to be inserted, let's do the test one more time to make sure nothing has happened.

Now, the truth is, end-to-end processing is way faster because we've automated things, a big variable back then, NTIA processing was a big variable, sometimes that would take some days, for example, like if you pass technical checks and everyone approves, like processing is very fast. So, supplemental arguably is vestigial and not needed anymore, but that's why it's there. I would suggest it still has a role for ccTLD transfer requests where, it could be months, many months, between initial submission and root zone implementation, but that's kind of a corner case.

So, our suggestion is that instead of trying to be too clever about like what is the first technical check versus what is a retest versus

what is supplemental, just consider every time we do a test run to be a test run in a just a general category for technical checks and we just measure end-to-end performance irrespective of the context of why we're doing the technical check and I think that would smooth out some of these reasons as to why we do the checks at certain phases of processing.

I mean, it'd still give useful measurements for the CSC to validate is this trending in the right direction, still flag potential issues, but that's a suggestion to consider. But yes.

PETER KOCH One question. Do I understand you correct that the test itself is exactly the same each time? So it's called differently, but what you test is exactly the same.

KIM DAVIES That's correct. Now in the context of an individual test run, what tests are performed varies depending on the nature of what you're asking to change. So, it's not like a very specific set of like 28 DNS queries or whatever you might say. It's like variable and dynamic based on the construct of your NS set and so forth. But yes, the fundamental business process is dispatched to a subsystem that just runs them the same way.

RICK WILHELM So, one, this is helpful, but it's kind of, as you sit here and think about these notions of the tech checks or the, the checks of the

TLDs DNS servers, when you think about the notion of an SLA measurement and such, can I take a step back and ask a kind of a silly macro question?

If you are IANA and your program, we'll just say the .Org name service, and you've got an SLA on your test. If my name servers are failing, why does that show up as a, and you're executing the test on time. If I'm not responsive, why does that show up as an SLA failure on your part?

Like, and I apologize if that is like the dumbest question ever, because I should know the answer to that. But like, in when I do SLA measurements, if like I do my part, and then if the client does not respond, I only have, because my SLA experience is like, which rooted and grew up in telecom management, where the concept of a DMARC and a boundary of my service interface, between me and customer premise equipment, and I apologize for all these telecom stuff for those who didn't, were not damaged in their youth by telecom, me and customer premise equipment, and that's where the line is. But if you're querying my name servers, and I'm not responding, why is that damaging your SLA measurement? Stupid question.

KIM DAVIES

No, I mean, I think you've spoken to the heart of the next point on the slide, which is, it's all good, it's all good. No, I mean, I think, again, that's at the heart of the question upon us. Like, I think I could argue both sides of the equation. So yes, like the time we're

waiting for packets to come back is not time that we really can control.

Well, I mean, we can set more aggressive timeouts, we can do less retries. I mean, there's ways of like, gaming it, gaming is probably the wrong word, but you know, tweaking it. But I wouldn't suggest that these measurements are valueless either. I think, we have meaningful dialogue in the CSC whenever there is a variance. And as we just did earlier today, we have a substandard discussion about what's the nature of happening. I think that's useful and constructive.

So, I don't want to pretend to know what the answer to this is. Obviously, one way is we could have with high level of precision measure, the tick tock of every packet we send in and out. And I mean, we can measure like in a very fine detail, all the different aspects of it. I don't know whether that's meaningful measurements for the purposes of what the CSC does. And I think that's an exploration we should do together to work out, is the juice worth the squeeze there? I mean, maybe if we look at all the different things here, maybe the current approach is the best out of a bad set of options. Because yes, it's not all time that IANA can reasonably control, yet it's the best thing we have to give CSC some visibility into trends on this aspect of our work.

And to be very clear, like particularly in the early years, 2016 to 2018, 2019, I would say that this measure yielded significant work from us to optimize what we were doing. Like, we had a highly

sequential process where, frankly, it wasn't particularly optimized. Like, it was technically correct the way we did the checks. But we didn't like reuse RR sets from previous queries that we could have and other things.

In the first few years, our coders like they started parallelizing operations so that they could be done. And it yielded meaningful improvements in processing time. But I feel like we've kind of yielded all those kinds of benefits that we can reasonably do at this point. And another fact to consider is that we're well overdue from doing some community consultation on what the technical checks actually are.

And we've heard a lot of people want us to do more. Like, there's a lot of other things we could be checking for that we don't today. I think that's a process that's a bit out of scope for here, but I think I raise it just to indicate that I don't think the scope of technical checks will be the same forever. And if we were to sort of materially change what we're checking, then that will obviously have an impact on these measures in some way.

PETER KOCH Yeah, thanks again for giving the mic to an observer. So, I will observe two things. One is we do technical checks for the second level as well. So, we've been through some of these discussions, even though we don't do them under SLA. However, the real observation is when the response doesn't arrive at IANA, it doesn't mean, Rick, that you didn't send it. There are network conditions

where that is a shared responsibility getting the response back to the tester. It's an edge case, most of the cases, but it is.

CLAUDIA RUIZ Please use the microphone.

RICK WILHELM Sorry, I also, I agree with Peter. Yeah, that's also true, share the joys of UDP and routing and whatnot.

KIM DAVIES All right, next slide, please.

RICK WILHELM A hand from Anil just went up.

KIM DAVIES Sure.

ANIL JAIN Yeah, thank you, Rick. I just want to interrupt the presentation at this point. We have discussed the non-technical SLAs and technical SLAs. Which are the major customer complaints which we are receiving, which may focus all of us to look at some modifications in the SLA. Thank you.

KIM DAVIES

I'll chime in from my perspective. I think we're doing this in the context of a lack of customer complaints, which is always a bit of a challenge and something we have across the board for IANA. We do customer surveys and the like, and generally speaking, our customers are quite happy.

So obviously, we want to continue to improve, but it's not in the context of there are any significant complaints about any of this. Broadly speaking, we're hitting all of our SLAs, as I think the CSC knows. I think that's a reflection of we're actually doing what our customers would like us to do, at least for now.

RICK WILHELM

I'll second that, and thank you, Anil, for the question.

ANIL JAIN

Thank you.

KIM DAVIES

I think the next major observation on the SLAs, and this certainly comes from my team. I'm a few steps removed from this these days, but one of our functions is processing label generation rule sets. So, these are the rules on how IDNs are registered within registries, basically, and one of IANA's responsibilities is maintaining a repository of, at this point, 10,000 plus tables, I think, and ccTLDs voluntarily contribute them, but gTLDs are contractually obligated to contribute them to IANA.

So posting gTLD LGRs is a big business for us, in a sense. That said, because it was a big business and it was a lot of work, certainly for when the last 2012 round of gTLDs was spinning up, we actually yielded a lot of process efficiencies by essentially optimizing the process where gTLD operators have to submit them twice, essentially.

They submit them to ICANN for approval on a contractual basis, and then they submit them to IANA for posting, also on a contractual basis, but a different contractual provision. So, since ICANN has already got them, we optimize that and essentially, since we know this is frankly, some gTLDs were forgetting to do the IANA piece, we optimize that and essentially, once ICANN has validated them, they pass them on internally, rather than putting that work on the customer to then submit them a second time.

So, this means that today, in terms of that normalized workflow, we're actually not getting them from gTLD managers anymore, we're getting them from another department within ICANN. So just we're measuring the wrong thing and whether we should measure it at all, I mean, there's an open question, but I think this is an opportunity to re-evaluate that business process, is the current measure fit for purpose?

RICK WILHELM Thank you, Kim. I actually don't pay much attention to the LGR SLA. Does it show up as a zero? Because coming from the actual customer of the TLDs, you get zero, but coming from the internal

customer, you get thousands, right? So, how does it show up in the SLAs?

KIM DAVIES

It still shows up as a request, but I think, and this comes back to why do we have SLAs. The SLAs, in my view, is that a customer submits a request to IANA, and in this instance, and how long does it take for the request to be completed? And it's basically fulfilling customers' expectations. But in the current process, it's all opaque because it's completely disconnected from the customer request.

Firstly, it's a contractual obligation. I imagine many gTLD operators are doing this out of obligation, not because they're desirous of seeing it pop up in the IANA repository, but because they submitted it in NSP, which is the name service portal, I don't know how long the ICANN side of it takes, but I believe it takes days or weeks or it varies.

And so, when it hits IANA, it's completely opaque to the customer. So, we're not really measuring something that is tangible or useful from the customer perspective of how that business process runs.

RICK WILHELM

Plus one to that. Gordon has a hand for each. Gordon.

GORDON DICK

Yeah, I think there's a couple of aspects in that when this was put in place, obviously, as a contractual requirement to have it

published to IANA, there was a need for the SLA to trigger now that it's on ICANN's responsibility to put it there. It's in their side of the contract rather than ours as a registry operator.

I think the other thing that makes me sit back and think there's change in this area more generally is the round two process where RSPs are pre-evaluated for particular IDN tables as well. Because what you're having to do at the moment is manage SLAs on a per TLD basis when actually in the future it seems to be more, yes, it'll be a TLD, but it's using particular table that's pre-evaluated from an RSP and there's no visibility of that connection and the equivalence at your end, which changes the dynamic of what the public record's there for, I suppose.

KIM DAVIES

Yeah, I mean, you make some good points, won't go too far down the rabbit hole in the interest of time, but even today we have some of that interesting complexity, like not to use names, but a large RSP that manages hundreds of TLDs and maybe makes changes to 10 LGR tables will submit it as a single request and say here's my list of TLDs and apply these new tables to every one of those hundred TLDs.

So, that is a significantly different level of work than just one TLD updating one table, yet it's still one request. So I think there are some challenges with the methodology even today, but you make a good point and that the use of reference LGRs, I think, set aside the SLAs for a moment. We're going to have to rethink our business

process for that as well. I mean, obviously there's efficiencies to be gained.

I think the question for this group is like, what is a performance standard that is important for the community in terms of what IANA does and what's there to measure and what's appropriate to report?

RICK WILHELM

Yes, completely agreed. And also with root zone changes, there's definitely temporal relevance to it, right? If you don't change this stuff, things start to break. When LGRs move at tectonic timeframes, right, they're very slow moving, right? And if it happens on Tuesday versus Wednesday, the world is not going to come to a screeching halt in most -- in 99.%, and keep adding 9% of the cases.

So yeah, it's a very different beast. That doesn't mean that IANA shouldn't be doing a good job. That doesn't mean these can't sit, it doesn't mean they should be allowed to sit around and it doesn't mean that IANA shouldn't strive for efficiency and all this stuff, but it's a very different thing than processing a name server change for a TLD and such. Thanks.

KIM DAVIES

Agreed. Next slide, please. I think probably won't elaborate too much on this right away, just to note that we are introducing a new business process for the Next Round of New gTLDs, which is

temporary delegation. This is a result of this NCAP process where essentially, we will put potential strings in the root zone for a short period of time, but it wouldn't be a delegation in a regular sense. Like we're not awarding responsibility to an entity to manage it.

It will be in a controlled manner, like very structured, systematized, and because it is so different from anything else we do today, we're not intending to like try and shoehorn it into what we're doing today. We're going to create it as a new business process.

We'll denote them differently in the root zone database as temporary and details to be worked out, but essentially, we won't like call them gTLDs, they'll just be something new. But what's important is they'll be highly automated, little stuff, interaction, et cetera. So, I think it's an open question of how we look at that, are SLAs merited or not, but I think one thing I will flag is this will for the first time test root zone growth limits. Like there are ceilings on how fast the root zone can grow, and this has a knock-on effect because the policies are such that when we hit the ceiling, like we have to stop adding TLDs at least until the next month. I think it's a month-to-month granularity.

So there's some exploration there, but I think I just note that from an SLA perspective, like we will probably want to look at some combination of monitoring TLD growth month to month, like percentage change in root zone size, things like that, in service of these new policies that will be implemented as part of the Next Round. What the actual burden on IANA will be in terms of like

commitments, unclear at this stage, but I think that's part of the discussion we have to have in the future.

RICK WILHELM

Hold for just a second. I've got a quick question. Go back to that one right there. Pablo has a hand. Real quick question, just a clarification. Right in the middle, related, it says related, there will be root zone growth. I don't follow the Next Round religiously. Is that a fact that there will be capital RZ growth limits, that there's going to be a number that IANA is going to be given, a speed governor, a speed limit that you can't exceed? 17 per day. I'm making the number up and just plucking one out of thin air.

KIM DAVIES

Exactly. It's a percentage growth limit and I should know this off the top of my head, but it's -- I mean, it will have a compounding effect. So, the limit today will be different from the limit in a year because it's a percentage of the current populace of the TLD, but there's some limit, and yes, if we hit that limit in any given month, we cannot add a new TLD and there'll be some level of management on our side and we've talked about our priority is going to be CCs over Gs for a start. Like we don't want to be in a position where if a new CC, as rare as it is, needs to be added, that we don't have room for it and we say too bad, so sad, wait till next month.

So obviously, that's part of the logic we have to have, but then also like what comes first, like temporary delegations over actual delegations and so there's some things to be worked out there, but fundamentally, yes, there'll be a ceiling and hopefully we manage the ceiling that we never have to like never have something queued up and they have to hit stop, but in that case we have to hit stop.

It's not because IANA is processing it, it's because IANA is obligated on policy not to proceed. So, I think we don't want to be penalized on our SLAs that IANA is taking a month to process this request, but we're not processing it, like we're obligated not to do it till the first of next month. So, that's kind of the situation.

RICK WILHELM: That's very helpful and good to get, let me go to Pablo really quick, and then Selina, you had a point.

PABLO RODRIGUEZ: Thanks, Rick. Kim, is these temporary delegations also a tool to act as a remedy on what happened with the situation with Lebanon?

KIM DAVIES: No, it's for a completely different purpose. This is about collecting intelligence on how non-existent TLDs are used and by virtue of putting them in the root zone it might, there might be aberrant behavior like software malfunctions or things like that that we weren't aware of until it was put in the root zone.

You know, things like there are like non-existent TLDs that are in like firmware of like home devices, things like that, and the notion is that once it's in the root zone things might happen that were unexpected. So it's a limited time window to monitor for the effects of adding a TLD that hasn't existed before and there's a whole policy around if something is observed then what that means for the viability of the TLD. But that's the purpose, it's not like caretaker delegations I think is what you're referring to.

PABLO RODRIGUEZ Thank you.

RICK WILHELM It's literally testing in production. You know the meme. I rarely test my code but when I do it's in production. We're literally testing the TLD in production. Selina, you had a point. Thank you.

SELINA HARRINGTON I was just going to say that the growth rate limit is 5% per month for the root zone.

RICK WILHELM Very good. Kim, please.

KIM DAVIES So, I think the next slide is just my concluding slide. I think you know main inherent flaw with the current SLAs is tech checks and

predominantly as we discuss many months of the year it's predominantly for reasons outside of our control.

Secondary flaw that we see as staff is there's all these categories, we went through a number of them today that either don't measure anything material or they used to but don't anymore.

New gTLD round will have an impact. We talked about name collision assessment et cetera, and that's a consideration. But the final concluding thought as I sort of touched on earlier in the presentation which is that I think we also have to balance the desire for a good set of SLAs in the modern time versus like the engineering costs. Like we invested a significant amount of effort to building the SLA regime we have today back in 2016 as part of the transition. That's a sunk cost.

Obviously, there's a maintenance burden and keeping it running which is not nothing but generally speaking it doesn't take a lot of effort to keep measuring what we're measuring. Significant changes to the SLA structure will require work, require planning, scheduling, prioritization and so forth. So, it doesn't come for free to make changes so I just leave you with that thought.

RICK WILHELM Very good, thank you, Kim. Pablo, Stale hand? That's okay, no problem. Thank you very helpful. I'm glad that we had good time to go through that. This is a good presentation. It's already posted

I think or it will be within shortly. Very good. Anybody have any concluding questions for Kim?

I think this sets us on a good path on the SLA review. I think that what we can do is, you and I can work together to take stuff out of your slide and put it into the proposed changes grid and we can start to some discussion at the next meeting or in between the meetings on those and kind of get those structured into the grid about proposed changes and then look at those and see if anybody else has some other ideas.

But I think that this presents a good point and I like the stuff that you had in the concluding slide about there is technical inertia in a positive way, like the code is in motion and like a spacecraft or a ship and if we want to change, it requires maintenance as it's in motion, but if we want to change its direction, that requires effort too. And so, we need to be conscious of how we expend our fuel to move the spaceship through space. The good ship IANA. Yes.

PABLO RODRIGUEZ Thank you. I want to take this opportunity to thank Kim and the rest of the team for the excellent work and presentation and as a comment. I used to in some of my messages and Selina, you may have seen this when I asked but last month, we had 100% this month we have 98%. What's going on? What changed?

And now I realize that there are things outside of IANA's control that is causing that some failures and we are considering that whatever

you're measuring fail but at the end of the day in reality it's not IANA's fault. So, I wonder if that is something to consider and to start measuring in a different way or something so that you don't need to own a failure that in reality is not really yours. Thanks.

RICK WILHELM

Thank you. Very good. I think are we at the next meetings. You can see those there. We don't meet at the June meeting at the B meeting as usual or no we yeah, no meeting and so we've got our meeting book there for 104. So those are just there and published.

One thing for AOB, Kim, did you want to make a quick comment about the RZMS and I didn't prep you for this I'm sorry but you make a quick comment about the RZMS logging thing the email that went out? Yeah, just briefly.

KIM DAVIES

Yes, certainly. So, as I think it was forwarded to the CSC list, right? So, we discovered that there was an issue with our resume management system. Essentially, it was logging information about credential usage in our internal system logs.

There's no indication that there was any exposure there or any usage of the credential data that was in the logs. We are doing an investigation to confirm our understanding and that's still ongoing but just out of an abundance of caution, we did a full reset of everyone's credentials. I think it was Saturday a weekend ago. Like

I said, this is us being very conservative and paranoid. Hopefully, there's nothing there.

I expect when we have a full understanding of exactly what happened, we'll give you more detail. But as of right now, this is a precaution. We asked everyone that has a resume management account to reset their password and I'm not sure there's much else I can really say at this point.

RICK WILHELM

That's fine. I just I wanted to make sure that we had in our meeting that we covered it since it was timely because it would have been a little bit irresponsible of me not to give you the opportunity to talk about it. Of course, many folks, many customers, have two-factor turned on their accounts, and if you've got two-factor turned on, that's obviously very, very good protection.

Not all CSC customers have two-factor turned on for reasons that have been discussed here before because as personnel shifts changes and things like that it can be problematic and dot, dot, dot and we don't need to kind of rehash that. I'm going to be giving a CSC update in the registry stakeholder group tomorrow. I'm going to use the opportunity to encourage the G's to be -- I think it's easier for G's and this is in my opinion. It's easier for G's to be having two-factor on their accounts typically.

So I'm going to make another call for G's to be having two-factor on their accounts. CC's can get a little bit more complicated and we'll

let the CC folks' figure that out for themselves, but obviously, if you've got two-factor turned on, password breaches are a lot less of a problem.

So, anybody else have any questions or comments about the topic that Kim just brought up there?

CLAUDIA RUIZ We need to wrap up.

RICK WILHELM We need to wrap up. We got two minutes though, right? Don't we go until 15 after the hour?

CLAUDIA RUIZ We had until 10:00.

RICK WILHELM Oh, I thought we had until 10:00. I thought the agenda said until 10:15.

CLAUDIA RUIZ Did it?

RICK WILHELM Agenda said 15, past.

CLAUDIA RUIZ	Okay.
RICK WILHELM	Oh, I'm sorry. I was frolicking along like we had until 15, until we had until 15 past.
CLAUDIA RUIZ	It's just when you said it for interpretation. Thank you.
RICK WILHELM	Okay, I'm very sorry. All right. Well, then we've got to stop. Good grief. We were adjourned. Thank you, everybody. I wondered why Peter laughed. All right. That's it, everybody. Thank you all very much for your time. Thank you very much, Kim and the rest of the PTI team. Everybody here and back home, we appreciate your help. Thank you, all. Cheers.

[END OF TRANSCRIPTION]