ICANN75 | AGM – Tech Day (4 of 4)
Monday, September 19 2022 – 16:30 to 17:30 KUL

EBERHARD LISSE: Welcome to the last block which is now being recorded. Eduardo Alvarez from ICANN will speak about his EAI survey tool with a focus on the technology behind it, rather than on the results. Please, you have the floor.

EDUARDO ALVAREZ: Thank you. Good afternoon, everyone. As Dr. Eberhard has said, my name is Eduardo Alvarez. I'm from the GDS.

EBERHARD LISSE: Can you take your mask off? It's easier for the people in the room to understand.

EDUARDO ALVAREZ: No problem. It's a pleasure actually. This is okay, right? Better. So I was saying part of the GDS team at ICANN and I'm here to talk about this new tool, not so new tool, that we implemented to measure the support of EAI email address internationalization on email servers.

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.

Next slide, please. So this is part of an initiative that ICANN org to launch as part of the work on universal acceptance with the purpose evaluating the status and the progress over time of UA readiness across email servers. The purpose of this tool is to measure the support on email exchanger servers on second domain names across TLDs. This tool uses as input the TLD zone files and the information available in the public DNS. And that's output, it will generate a list of classified on part TLD country and the support rate for the SLDs, the MX servers found, and the IP address that those MX servers resolve too.

For the process of the survey, can we go to the next slide, please? I guess it can be categorized in these three steps. First, taking the input of the TLD zone files, the tool will pass all of the second level domain names in the zone file, and then we'll query the public DNS for MX records under those second level domain names and resolve those named servers to their IP addresses. Once that full list is compiled, it goes to step number two, which is actually performing the survey tests. It will retrieve the IP geolocalization details using a third party library, and they will proceed to get each of those IP addresses, the MX record resolve to perform a test with two email addresses that use internationalized domain names.

First, as we see in the screen, it will use a test address confirmed by non-ASCII, name of the email address, then, domain name in the form of ASCII label, and then it will execute the trigger marks that we see in there. First, it will initiate a SMTP transaction with extend hello command to the server, and look for the appropriate header, which is SMTPUTF8. This is described in RFC 6531 for the internationalized email address usage. Once if the header is found in the response from the server, it will initiate an email transaction issuing a mail from command with the test addresses that I described before, and then it will look for a positive response, which means code 250 and then it will terminate the connection.

It will repeat the same process again, once for each of these two email addresses and record the result. So then on to the last step, the tool will just store the results for that IP addresses that were just tested, and it will aggregate it also in terms of the mail exchanger server. Classifying it does full support if all of the IP addresses that the mail server resolved to, pass the test. Partial, if some IP addresses passed the test and some didn't, or no support if all of the IP addresses failed the tests.

For an example, can we go to the next slide, please? So this would be sort of how expected transaction would look like to pass the test of proper EAI support. So it will first start a connection with the IP addresses of the email exchanger server, and it will start with the extended Hello command, which we see highlighted in blue there. So it will identify itself, and then the server is expected to respond with the proper headers. We see it highlighted in red, the one we're looking for is the SMTPUTF8 as described in the RFC. And then if that is found in the response, then it moves on to the next step. If it's not, then it's already recorded that it’s not EAI compliant, and stops the test right there.

But if the header is present and it proceeds to initiate the mail from command using those internationalized addresses, as we see in this example, it's known as key owner name and then the u-label of the domain name used for testing. The second variant of that email address as I mentioned before, it would be the same owner name or name of the email addresses, but the domain name would be in the form of an A label. So based on the response to that command, it will look for the positive response, which is what we see there also highlighted in red a 250 response code, and then it will terminate the connection.

Can we go to the next slide, please? So that's roughly the process it does to do the test for each of the servers. Now on to the tool details. This tool was implemented in Java. It's compatible with Java version 8 or later. It also requires Apache Maven to build the tool. To run it requires the runtime environment for Java, same version 8 or later. It requires Docker, MariaDB database management system to store the results locally. Our implementation also supports uploading the results to cloud a data warehouse solution which is a Snowflake.

For that it uses a client known as SnowSQL and the user that runs the tool must obviously have proper permissions to run Docker and be able to run Sudo to run the commands in the server. For more details, the documentation, configuration, description, and full source code is available at ICANN's Github repository. We have the address in the screen that's github.com/icann/eai-survey-tool separated by hyphens. So that's available for everyone who's interested in looking in more detail.

Can we go to the next slide, please? So to go a little bit more on the data model that these tool uses, here's the database schemata that the tool requires. I'll go from left to right describing what data we're storing here. So, first on the progress table, this is just to measure the times where each of the steps of testing is executed and completed. And then, from the other tables at the right, you'll see that's where the actual data from the second level domain names, the MX record found, the IP addresses that those MX records are sold to are stored. And it's normalized in a way that every IP address is stored only once, tested only once, regardless of how many second level domains may reference the same MX servers or how many MX servers may result to the same IP addresses.

Can we go to the next slide, please? So in terms of third party services used by the tool, it's configured to run-in AWS virtual environment. So that keeps our survey outside of ICANN'S internal infrastructure. So it runs like it's a public server with normal access. It also relies on Google's public resolver to resolve the second level domains names for the MX records in the DNS. For geolocation services, as I mentioned before, the third-party service we use is from a company called Maxmind, which is a database that just keeps among other stuff the relationship between IP addresses and the country that those are registered on.

And lastly, Snowflake, which I mentioned briefly before, that's our service used to persist our results in the cloud, this is data warehousing service. Just to clarify exporting these results as snowflake and the cloud is something that we do at ICANN, but if someone decides to check out the tool and run it on their own, that's sort of an optional step that can be skipped.

Next slide, please. So with our methodology, there are some risks of false negatives, and we've seen it a little bit in our results. So we had to implement a few mitigation measures to reduce these instances. The server where we run the survey, which as I mentioned before, is an AWS instance, has a functional SMTP server, so the mail server is running and is able to send and receive email for the test addresses that we're using with different scripts, and for each survey that we run, it uses load balancing with sets of two IPv4 addresses and two IPv6 addresses since we're testing millions of domain names and MX records.

We also switched to a different set of IP addresses to avoid or to verify if there was some impact in terms of some sort of restriction or blacklisting occurring due to possible classification of spam servers, even though as I mentioned a little bit earlier there, our survey tool does not send any emails. It just initiates the transaction, but does not actually send any email to avoid or reduce the risk of being flagged as spam. Also, there is really no need to send the emails. We don't verify the actual receipt of any email. The IP addresses that we're testing are only tested once, as I mentioned before, no matter how many times those are referenced by different second level domain names or MX servers. It also implements a maximum number of retraced.

Currently, we just try one more if we don't get a positive response from the server and we also implemented a waiting time of 30 seconds between each test, for this I might be, which is only once. This will also help avoid any issue where our tools consecutively tries to initiate transactions with MX servers and be classified as some sort of attack. And not only between retries within the same IP addresses, but we also implemented the mechanism that it will have some wait times between testing while they might be addresses within the same /16 IPv4 block or the same /48 IPv6 block since those are typically maybe in the same infrastructure or same company.

Can we go to the next site, please? So as I was saying before, there are some known issues mainly or summarize here, we don't have a way to avoid running into automatic DNS block lists or as also called anti-spam traps where these are MX servers that are publishing the DNS. But if someone initiates a transaction like our survey does, this will be immediately flagged as a spammer, since it's not expected to receive an email from anyone. Also, we have also seen email servers that may have strict security policies. So if they're receiving contact from unknown clients like our survey tool, then that's immediately considered as unsolicited traffic, and they can also flag us as malicious or just spam basically as a bad actor. That's roughly the description of the tool. I want to go really quickly over some of the results that we've seen.

Can we go to the next slide? Some of the statistics that auto-generates and all of this is aggregated. We run our survey quarterly. We've been running it for the past year. I'm just showing that last three quarters to illustrate each time our survey run, which is each of the columns in this slide, we use a different script just to also mitigate the fact that some scripts may have better support than others, which as we've seen, that does not seem to have any impact. Right now, we've run the survey only using gTLD zones, which is what ICANN has access to, and we can see how many TLDs are available at the time of the survey in the corresponding row below.

The number of unique MX servers found across all of the second level domain names under the gTLDs that ICANN has access to the zone's file for, we see there is roughly space relatively stable around 35 million but we see a small decline in the numbers. In the last survey from the past month of July. And the same for the unique IP addresses that all of these MX servers resolve too. We see a trend of a slight decline in the number of IP addresses that all of the MX servers resolve too. And just for clarification or for reference, we also see that these IP addresses are 94% IPv4 addresses and only 6% are IPv6.

So, the results that our tests have seen, we see that the EAI support is relatively low if we're looking at unique IP addresses. So it stays roughly about 7.24, ranging all the way up to 7.38. In the last survey, we see a small pattern of the increase in the percentage there. But we still have 59% approximately of tests that are not passing whether the servers are not responding with the required header or when we initiate the transaction with the mail from, we're not receiving the positive response we might receive some negative response that is not supported, that is malformed or some other response.

And lastly, we also see these 33% approximately of IP address that could not be tested, which could be either because the server did not respond, refused the connection from our client, or responded with some other error code in SMTP, for example, 500, which are error or 400, which may have been other policies like the connection was refused or that the server was overloaded. And we also see IP addresses that are private IP addresses. So those are skipped by our tool. For some reason, they're published like that in the DNS, but those fall into this category as well.

Can we go to the next slide, please? Just to show up quick contrast, this is the same chart from before, but focusing on the values per mail exchange server. So, while we still have that 7% before for IP addresses, if we look at the actual MX records, the percentage is much higher. It's goes from 19 and on our last survey all the way up to almost 21% of full support for MX servers, which means that this almost 21% of the MX servers are resolving to some of those seven point something percent of IP addresses that are actually compliant for all of their IP addresses.

It's a little bit more optimistic in this approach, but we can still see that 67% which went all the way down to 65% in the last quarter of mail exchange service that do not offer any support. And then we see just a little bit more broken down the IP address that could not be tested or MX servers that did not result to any IP addresses. We're also separating those there. Which also I think it's a bit interesting to see in the tool.

Can we go to the next slide, please? We probably don't need to go too much into detail, but this is just less aggregate data in total, but more on a per TLD basis. This is the kind of information that the tool generates. Per each TLD zone file, we can capture how many second level domain names were there, the MX records found for those second level domain names, and then the characterization I described earlier on whether they have full support, partial support, or no support based on the each of the IP addresses that those MX records resolve to.

Can we go to the next slide, please? Another visualization from those results can be per country code since that's one of our third-party services that we're using to assign a country to each of these IP addresses that were tested. Again, this may be helpful for some. Or depending on how you selected data, it can be more meaningful, but it's also available by the tool. Obviously, countries with fewer IP addresses might reflect higher rates. This is a percentage not total. So that's something to keep in mind.

Next slide, please. What's next for our survey tool? There are some things in store for us that we're working with versus integration with the (ITHI) Identified Technology Health Indicators project, which is run by our colleagues from the office of the CTO at ICAAN to publish some of these metrics. We're also working on implementing some of these changes here, identifying the mail server software based on the test results to more easily identify what provider can be having this issue and see if there are any patterns or trends there.

We're also working on implementing to the change to persist if this additional header 8BITMIME is present or not in the response. This is also mentioned in the RFC for email address internalization, but it's currently not being verified. And some other internal changes to Docker as application make more flexible configuration to disable the geolocation feature for those that are not interested in that. Some internal notifications.

Next slide, please. Next slide. Thanks. And some support for other RFCs, support for SMTP over TLS. Configuration of DKIM in our test server. DEMARC as well. STARTTLS. Null MX. And some logic for better handling of errors when doing our test just to clean up a little bit our measurements, but that's also in our headline of changes.

Next slide, please, up, I think go left the presentation. Can we get the slides back? Well, I'll just hurry. There we go. Thank you very much. So it is the last slide anyways. So as I mentioned before, this survey has been run for gTLDs only. But I wanted to also mention ccTLDs are welcome to participate if they want to get measurements on these support rate for email address internationalization.

I agree, depending on these regions, it can be more valuable to know whether other scripts and languages are supported properly. So there's the option for anyone interested. You can just send us an email to GlobalSupport@icann.org, requesting more information, and we can help you. We can include the TLD in our survey if you want to, or you can check out the tool on your own if you prefer.

That's another option. But if you decide to have ICANN grant the survey for you, there's also other opportunities to participate or other tools that we offer that are also available, like participating in the DAAR project that also the office of the CTL runs or access to our monitoring system API for statistic for life statistics on the critical service availability of DNS or WHOIS. We do have about 32 ccTLds already participating in the monitoring system API and 21 in DAAR. So we're always welcoming more people that are interested. And that's it for me. Next slide, please. Just questions.

EBERHARD LISSE: Thank you very much. Can one get the source code of this? I think I would be quite interested in maybe adapting a few tools to test my own.

EDUARDO ALVAREZ: Absolutely. The source code is available in the GitHub repository I mentioned before. It's in the slides. That's github.com/icann/eai-survey-tool.

EBERHARD LISSE: Excellent. Any questions from the floor?

NABIL BENAMAR: Yes, this is Nabil Benamar, from ESG. I'd like to ask you about the statistics related to IPv4 and IPv6. So you say that only 6% of the tests were done using IPv6 which doesn't reflect the current state of the traffic Internet. More than 40% traffic is on IPv6. How can you explain this?

EDUARDO ALVAREZ: Well, thank you for the question. So what we see is when resolving the MX records found for under SLDs and gTLDs, that was the result we saw. 94% of the IP address those resolve were IPv4 and only 6% were IPv6. I mean, we do have the data. And this can be reproduced as well. If you run the tool, you can also get those results as well, and hopefully confirm.

NABIL BENAMAR: So what does this mean? Does this mean that the server is not the deployed on dual stack or something?

EDUARDO ALVAREZ: So it means in the DNS, the MX record did not have an A record or a quad A record.

NABIL BENAMAR: Okay.

EBERHARD LISSE: Okay. We are running a little bit over time, so I'm going to stop this here. Thank you very much. Thank you.

EDUARDO ALVAREZ: Thank you.

EBERHARD LISSE: So now we are coming to what I call the main event. Kim Davies is going to give us some insights into the current thinking and plan for the root zone management system.

KIM DAVIES: Thanks, Eberhard. Hi, everyone. I hope I fulfill expectations. So this presentation is really about an effort that's been underway within our team for a number of years now to evolve the root zone management system. So without further ado, next slide, please.

So the crux of why we're here today is that later this year we plan to introduce a next generation root zone management system. The root zone management system, which I'll describe in a moment, is one of the key platforms that we use within the IANA to deliver our services. And this next generation platform will introduce some important new changes that we think will benefit our customers and also set us up in good stead for future evolution for root zone management.

So today, I'm going to talk a little bit about what's in this next release, but also what we have in our thinking beyond that down the road. To be clear, this release that we're playing later this year is the first of what we hope to be many. So it's not the complete delivery of everything we've been hoping to do, but a first step in that direction. Since it's related, I also wanted to use the opportunity today given your technical audience to talk about some of the other technical aspects of root zone management that we're looking to evolve in the future.

Next slide, please. So for those that aren't familiar with the root zone management system, what is RZMS? RZMS is a system that we implemented about 12 years ago. And it manages the workflow of most roots zone change requests from the moment that a TLD manager submits them to IANA, through their processing, all the way through to implementation. You might have heard I refer to as root zone automation. That's a term I try to avoid.

Whilst automating a lot of the workflow is an essential component of what RZMS does, it doesn't fully automate the root zone. It shepherds requests through all the different processing stages, where there's an opportunity for automation, for example, technical check performance that is automated, things like interacting, sending emails from the system, notifying you have updates, and so forth.

That's automated. But not all elements of root zone management are automatable at least under the current policy. So our team is involved at certain phases of requests in processing them. But the system helps with those. It tells our staff when requests are in the right phase to do certain kinds of processing. So it definitely adds value in that regard.

A key part of the root zone management system is a self-service portal. For those of you in the room that are TLD managers, you'll be familiar with this. Using your username and password, you can log in, submit change requests. If we have interventions throughout processing of a request, those can be made in the system to a certain extent. You can check the status of a pending request. You can look at the history of your request and so forth.

The system integrates with other systems most notably is the root zone maintainer system. Verisign has a role as one of our root zone partners to essentially publish and distribute the finalized root zone file to the root server operators. And to do that, we send EPP commands to Verisign that represent the deltas to the root zone. Once we've processed change requests all the way through, satisfied ourselves that they're ready to be inserted in the root zone, we affect that by sending those deltas via EPP to Verisign.

Another integration we have is ICANN's name server portal for contracted parties, gTLD operators in particular. This allows for some integrations such as when you're establishing your new gTLD, there's a handoff from that portal to RZMS so that you can continue the process through the delegation in a seamless manner. And that's something that we'll be exploring and looking to possibly expand or reevaluate with the next round.

Now, RZMS itself, whilst I said it was launched in 2010, actually, it dates back earlier. If you look at the lineage of RZMS, it comes back to a project, kind of it came out of CENTR. Declaration, I was working at CENTR back then. So CENTR was exploring ideas for root zone automation as an outside party, and commissioned NASK, the Polish top level domain registry to create an experimental proof-of-concept. And that proof-of-concept is actually what ICANN build this platform on that we use today. So it goes back some 20 years.

Next slide, please. So what's the need for change? Well, the platform has grown over the years. But we've assessed that it's really constrained in terms of its architecture and the way it's built from supporting future needs. Why is this? Well, let's think about what this community was like 20 years ago. Almost every TLD manager with a couple of exceptions only operated one TLD, so one TLD per operator. No one was using smartphones to access website back then, maybe WAP was something you might have used as a curiosity, but there was no need for that. Obviously, the IANA functions were under the encumbrances of the NTIA relationship, which included every recent change had to go through NTIA review before being implemented.

And then the fundamental architecture, the software architecture, and the frameworks used to implement the system, date back to the early 2000s. So they're not certainly not contemporary today by any means. And also, the need for changes driven by pain points that we've observed over the last 10 and more years from our customers and also from our staff. A lot of what happens in root zone management today was really informed by that original contact model used by internet back in the 1990s. Admin and tech contact is a prime example. But what we've seen is that customers have evolved, increasing use of role accounts. It's not so much people, but roles within organization.

And we've also seen that that model straining with requirements, which has led to a lot of manual interactions where customers with specific needs are often working directly with our staff, rather than using the automation system. It's actually easier to reach out to our staff, tell us manually what they're trying to achieve, and then we will instrument that in the back end. Another factor is that public points of contacts are spam magnet. That their public information, it's in WHOIS, and that can be problematic.

Coming back to some of those complex operational requirements, of them, the most obvious one is bulk updates. So we have some of our customers that operate 200+ TLDs and some of the kinds of changes they might want to make need to be reflected in every single one of those.

The current model, again, coming back to the one domain per TLD, necessitates you to submit 200 change requests to change one thing across 200 TLDs, which is obviously not optimal. And then lastly under the pre-transition environment, IANA was certainly constrained in exploring different modes of operation, but that restriction has been lifted. So we now have more flexibility to explore improvements in these areas.

Next slide, please. So what's new? So we're working to deploy in the coming months, is comprised of the following. Firstly, it's a complete platform rewrite. It's been rebuilt from scratch with modern architecture. We discussed what we wanted to accomplish with this system with ICANN's engineering and IT department, and their advice was essentially to raise it to the ground and let's build it again. So that's what happened. A key deployment in the first phase will be a new technical check system. Now this doesn't change the technical check methodology, but it does change the architecture behind the scenes.

So today RZMS is a monolithic application, and when you do a technical check, it's done within the context of RZMS. This has posed scaling challenges. It's not easy in the current environment to do parallel tests, scale it to cater for higher loads and so forth. So this should help address that by having it as an independent microservice that RZMS calls out to, to perform the technical checks. What it also lets us do is develop that technical check system in its own cadence rather than if we wanted to tweak or adjust the technical check how it's performed, we don't have to go back and modify the core system, but we can do it separately in the technical check system.

Now we've built it to generate comprehensive debug logging style logs. One of the criticisms of the current system is there is a technical check issue. It's not necessarily evident exactly what's happening. The current system is very brief in its descriptions. But here, you'll be able to download something very similar to a syslog with precise time stamps, exactly what was sent. Packet captures those kinds of things so you can really drill down into specifics. And we'll also add richer explanations. So more customer friendly text in the UI. So we'll actually explain with a bit more detail what's going on and what's been identified.

Next slide, please. Another modification that's pretty fundamental is our authorization model. Today we have admin and tech contacts and they serve two roles. One is admin and tech contacts are listed in WHOIS as points of contact for top level domains, but they also serve a role as cross authorizing changes within IANA. This can be a problem where for business reasons at TLD, those roles are split.

The people that act as customer service to the community that might be responding to inquiries that come by WHOIS records are different from the people you want to authorize, change request, which might be very fundamental in nature, like changing your contacts or ultimately transferring your TLD. So these are pretty important things to get validated correctly.

So what we've essentially done is split those two responsibilities into two separate data models. We retain admin and tech contacts, but they become just public points of contact. They'll still be in RDAP and WHOIS, but they won't necessarily be required to cross authorize changes. Instead, TLD managers will be able to create users in the system, and it can be two users that map to the existing admin and tech contact.

And, indeed, as we transition, that's exactly what's going to happen by default. But you will have the ability to go in and create new users. You can add a third or fourth or fifth person. You can just have one. You can have dozens if you like. And you'll be able to configure their access rights. So not every user has automatic access to do everything. You can create users that have limited access.

So this should enable you to support a lot of common ask. So for example, the TLD manager might have an RSP that's providing name service for them. And they might want to give the RSP access just to alter NS and DS records only and not touch anything else. So that will be possible in this new model.

Another change part of that is these users will be individuals. They will not be role accounts. So part of enhancing the security is we want each individual to have their own unique credentials so that when people move on from organizations, there's not shared credentials, which will be an improvement, but it will also help us, and I'll get to this in a moment, identify who's behind an account. So if there's ever a need to restore access, that can often be difficult with role accounts. But if we know a person's name, that is easier.

Another development I wanted to flag was shared glue improvements. Today, this can be a sticking point. If you change shared glue that shared amongst many TLDs, we require that the context for all these TLDs currently consent, which can be time consuming and laborious. So the new model is relatively simple. We receive a glue change. It only has to be approved by the TLD that requested it. All other impacted TLDs will be notified and they'll be given a 14 day window to object. So rather than being and opt-in, it's more of an opt-out model. If we don't hear any objections, which I think, frankly, we never do, 14 days later, the request will proceed.

Next slide, please. So I talked about the changes that you're immediately going to see with this next deployment. But we do have other things in the pipeline. In the interest of getting this launched, we deferred some of the functionality, and I'll talk about the key things. One of them is API access. This will give you a programmatic ability to lodge a request, interact with them, do much of the capabilities that are in the UI.

Really, we're targeting this initial API release at again those bulk users. Parties that operate tens, if not hundreds of TLDs and they have the need to do regular key rollovers or maybe one of their staff members has left, so they need to update the admin contact across how does the TLDs, stuff like that. So we think that while it doesn't solve all the problems, the API will greatly improve the way we conduct those kinds of operations. So the API will be very standard modern model JSON, HTTP endpoint, you'll issue revocable tokens, and connect to it with that.

Another improvement is in how we do technical checks. Today, it's a pass fail model. You either pass each individual test or you fail them. And if you fail them, that requires you to start talking with IANA staff to either convince us that it's actually not a legitimate failure or there'll be a dialogue and you need to go correct the issue. We will add a third category, which is warnings. And the purpose of warnings will be we might flag issues that you can self-dismiss. You can review it and say, "Okay, I understand, but fine, I don't care, go ahead." That will be possible in the UI without any interaction with staff. So we think that will provide a big benefit there.

Next slide, please. The big one is multifactor authentication. So I could talk for an hour on this topic alone, but current state of affairs is some of our customers have definitely asked for it, but we have conflicting advice. One piece of advice came from

ICANN's security and stability review team, which urged us and recommended that we implemented.

But we also had a recently published root zone update study that derives from the IANA transition that actually recommended that we don't implement it for complicated reasons. But it's suffice it to say, there's not a clear consensus opinion on this matter. So we're going to implement it as an opt-in option for those that want to use it, at least for now. Part of the concerns and part of why it's not a slam dunk to implement this is the low rate of interaction our customers have.

A lot of TLDs come to us like every five or ten years. And the likelihood that they've lost their credential or that staff member no longer works there, is high. And multifactor authentication is not very useful if you just go, well, okay, then we'll just reset your account like the multifactor was never there. So we need to develop an operational model that works well, that we can reboot trust, it's not a technical problem to put it mildly. It's an operational problem.

And we need to develop a way of having more robust know your customer procedures so that we can confidently reboot trust as we implement multifactor authentication. And I'll also note that we literally have customers in every single country in the world, including countries with sanctions. And so some of the options are not available to us for countries with sanctions. So that's something we need to keep in mind.

We also want to limit third party dependencies. So there are commercial solutions to some of this stuff, but we need a root zone management system that works when the Internet is not working. So limiting our dependencies on third parties for critical infrastructure for us is important. And for us, this means things like TOTP, web authentication, not these third partied solutions, single sign on solutions, and certainly not depending on the cell phone network for that.

So yeah. And then another consideration that I think is important to flag is a lot of our current model of trust with our customers is based on the fact that you control your top level domain zone. So by inserting records in your child zone, that is a pretty compelling evidence that the request you're submitting for the root zone is actually what you want, like, we checked that the DNS records at the child are being updated before you asked them to be put in the root zone.

That shows that you have access to edit the zone. And you already essentially have the keys to the kingdom at that point. So that's important. DS records as well. We ask that DNS key records are updated in the child before we put the DS record in the parent for the same reason. So you're proving that you already have fundamental control of the registry by being able to do that.

Next slide, please. So I also wanted to talk briefly just to finish up some of the other sort of technical related evolution that we're thinking about. Obviously, there'll be impacts on RZMS at some point, but it's not necessarily directly related to that. First is evolving the technical checks that we actually perform. The technical checks that we do today, the battery of tests that we run against TLDs, is really informed by a public comment consultation we did in 2007. A lot has changed in the last 15 years, and we think it's time to reevaluate that. The root zone update study I mentioned before has already provided some suggestions on how we can evolve that.

Privately, from a lot of our customers, we've being given advice, recommendations, or sometimes yelled at us about how we should update things. And I've been personally quietly writing that down and maintaining a record of that. I think what's going to happen is we're going to turn that into something of a white paper or discussion paper that will then go into a consultation and a discussion with the community on how to evolve it.

And I think importantly, coming back to RZMS, once we have that pass fail one system in place, I think we'll be in a good position to be able to implement a lot of additional tests that are discretionary because they can just be warnings. So we can flag potential issues and knowing they won't block request progress, but we have the opportunity to flag that to the customer. I think it will be a positive development.

Next slide, please. Another key thing that's related to this and also came out of the recent update study is proactive testing. Today, we only test when we receive a change request. This means that if a TLD falls out of let's say, compliance over time, there's some new emerging issue with their operation, they probably won't hear about it from us. They might hear about it from others, but not from us.

Proactive testing would be us just regularly routinely monitoring all TLDs for these factors. And if we notice some kind of change, we can notify the TLD manager. They might dismiss the issue, but it might also trigger a change request in the root zone. So, oh, yeah, I did update that NS record. I'd better update it in the root zone, for example.

Part and parcel of this is implementing things like CDS, ceasing these kinds of things. Like once we're monitoring proactively, we can look for signals that you generate to update records with us, and that might trigger creation of a change request. And part of this is, of course, we'll give you the ability to suppress these notifications. If you don't want them, you don't have to get them and so forth.

But I think riffing off this, I think a logical extension of this is some kind of health check panel in the root zone management system where at any time you can just review how your TLD is performing. And that wouldn't necessarily just be against technical checks. I could foresee and this is just me ideating, it's not a commitment, but we can report you have an old password, you might want to update it.

I know the TOD ops community talks about vulnerability alerts, and maybe there's some opportunity to flag those in the interface. Maybe periodically, we can validate your contact methods work in so forth. So these are all ideas to explore, but I think that they're a positive way of looking how we can evolve the service.

Next slide, please. So some of these elements we're kicking off engagement on at the forthcoming ICANN DNS Symposium. This was announced I think last week. It will be held in Brussels in November right after the ITF meeting. And so we're doing something new. We're having, I think, we settled on the name IANA Community Day, and because it's going to be a half day focus just on these topics, the technical discussion about evolving our technical checks. And also, another thing I didn't talk about today, which is an algorithm rollover in the DNS root zone.

So if you're in the area, if you're going to the IDS, I would encourage you to come to that session. But I know not everyone can come. And certainly, our intention is not just focus on that. This will just be an opportunity to kick off some of these efforts. But we'll, for certain, be coming back to future ICANN meetings and other forums to talk about this. There'll be public comment periods, no doubt. There'll be plenty of opportunity for engagement on these topics.

But for those interested, for those who have contributions, that'll be very welcome. But your thoughts are welcome at any time. So if you have thoughts on this it's welcome today, it's welcome this week. Drop me an email. It's always welcome. So with that, I think that might be my last slide. Let's see. Yes. Here we go. So thank you very much.

EBERHARD LISSE: Thank you very much. I for one will travel to Brussels. Any questions from the floor? Identify yourself to the colleagues.

MICHAEL PALAGE: Michael Palage. So, Kim, excellent presentation. Looking forward to seeing this rollout in the near future. The one question that popped into my mind and maybe you can go to Göran and ask him this question. So the SSAD light system they decided to build that on the existing centralized zone file access system. It seems like what you have here is so much more robust and has all the features that was originally envisioned by SSAD and it just seems like ICANN choose to build on a Ford Pinto as opposed to a Lexus. It would really be helpful if perhaps this Lexus could be used for the SSAD light.

KIM DAVIES: So I will punt on that. I mean, I don't know enough about SSAD to comment intelligently on that. I will say that, I mean, it is a specialized platform for root zone management. It's not a generic platform, so I don't know how true that is, but I'll definitely take that feedback back.

EBERHARD LISSE: Warren.

WARREN KUMARI: Thank you, Warren Kumari. So apologies if you covered this, but I'm assuming that you will still always have a way to make emergency changes. Right? Like you, as IANA, will be able to if the world goes, boom, fix whatever. Great. I assume so.

KIM DAVIES: Yeah. The capability to do emergency changes is not changed. It's the same. I think it's a backlog item, but definitely, we've talked about actually enhancing that. Today, to declare an emergency, you submit a change request, but then you call a call center. But there's actually a potential there that in the submission process, if you have Internet access, obviously, some emergencies you don't. But if you do, being able to actually flag it as an emergency in band as part of the submission process, and we could trigger all sorts of clever things. I mean, it's just an idea at the moment, but I think that that's a potential evolution that we could say. Thanks.

EBERHARD LISSE: Don't worry, my African brothers-in-law, we will always use your emergency contact. Elvin Lansing.

ELVIN LANSING: Elvinlansing.uk. Thanks, Kim. I love it. Thank you. Just quick question. The community today, is there online participation for that?

KIM DAVIES: I have to assume yes. I I'm not a 100% sure, but I believe that to be true. But you do remind me, I think I skipped over in the slides. We are having a session on ZRMS, including a demo. It's on the agenda tomorrow. It will not be technically focused, more customer service focused in nature, but if you're interested in more detail, please come along to the session tomorrow.

EBERHARD LISSE: Peter Koch.

PETER KOCH: Yeah. Thanks, Koch. This is Peter Lowe from DENIC for the record. Great stuff, Kim. I really love the references back to the NASK system and also shared glue is one of my favorites. Great to see this addressed. Let's now get rid of the discrimination against the WDS, and we're all fine. But on a more serious side, you envision that this is going to be an incremental set of changes. And that might sound interesting and valuable.

Many of us are running or categorized as critical infrastructure. And this is a very crucial change to the system because, of course, the delegation from the rule is in our critical path, in the risk assessment and everything around that. So when can we expect tangible written documentation so that we can evaluate that on the basis or to make security assessment and also develop internal changes to the extent that we need to adjust our own processes.

KIM DAVIES: Thanks. No, it's a very good point. I think generally, we see an opportunity to greatly enhance that documentation generally, but also specifically to the root zone management system. So that's something that we're working on, but it's a good reminder that this needs to be an area of focus. It's actually again that root zone update study, that's another one of the findings.

Without belaboring the past, there were constraints on publishing documentation, and I think it's taken a while to shake us out of that mindset, but looking forward, we expect to have much fuller documentation on a lot of these things. But if you have specific ask just to make sure we hit those notes, then please let me know the specifics as well.

EBERHARD LISSE: Any remote questions?

KATHY SCHNITT: Yes, Eberhard. Question for Kim. Have you had any plans to make beta testing the system or any other public available for registries testing opportunity before the official releasing? If it is so, how registries can participate.

KIM DAVIES: Thanks. We don't have a public beta testing program for this. We've certainly had internal testing and we've piloted ideas with certain customers and discussions and so forth to make sure that it fit their needs. That's not something we've contemplated at least not for this release, I think. But also particularly when it comes to API access, we're cognizant of things like having a sandbox and things like that so that in order to test your integration with RZMS, you don't have to first shoot on the production environment. Thanks.

KATHY SCHNITT: Eberhard, we had one more.

EBERHARD LISSE: Carry on with the remote questions.

KATHY SCHNITT: Thank you. Recognizing this is root level, but is there a way TLD admins can identify their TLD subspaces and generate these as a list.

KIM DAVIES: I'm trying to think of where this question is heading is that to do it public suffixes, or?

KATHY SCHNITT: They have in parentheses ieco.uk or com.au within UK, or .au respectively.

KIM DAVIES: I mean, we certainly have no plans for that right now, but I think I'll just make the general observation that we're here to service the needs of the community and so far as we can add functionality that is purposeful and aligned with what we're here to do. We're happy to explore it. I mean, it sounds like the nature of that query to me connects to the public suffix list, which we've had discussions on and off over the years.

Can IANA support the public suffix list in a meaningful way? Having TLD managers, for example, declare those public suffixes within their bailiwick and we could export that in some useful fashion for other parties to use. I don't know if that's a good idea or not, but that's certainly something that's being considered.

But on another topic I know in the TLD ops group at the last ICANN meeting there was a notion of should IANA add a security contact field to RZMS, which could then be exported to that efforts. So I think these are just two examples of value adds that IANA is well positioned to support because we have the trust relationship with all the TLDs. We have no active plans to launch them certainly not in the near term for what we're talking about today, but we're always receptive to these ideas. And based on the priority, we will factor them into future planning based on need and cost and all those fun things.

KATHY SCHNITT: Eberhard, that's it remotely.

EBERHARD LISSE: Okay. Thank you very much. Anything from the floor now? That leaves me to thank Kim. He didn't disappoint. Okay. And Brett Carr is then going to do the usual wrap up.

BRETT CARR: Thank you, Eberhard. I'm conscious in the way of anybody whatever plans they've got next, I'll try not to be too verbose. So a quick summary of what we've been through today. We started the day off with Charlene taking us through various case studies around IoT, and she showed the differences between 4G and 5G IoT devices. It was an interesting presentation though, but my knowledge of mobile networks is not good. So it was above me a little. I was a little disappointed about the IoT position on DNSSEC. It looks like there's some work to do there.

We then had Peter Thomassen in presenting an innovative approach from deSEC and automoted DNSSEC bootstrapping. It is definitely something to keep an eye on here as I think he's likely to become a standard in the ITF soon. Obviously, he's been supported by some big plays already as well, so it's good. Next, Michel and Marc reported on some great work related to universal acceptance on registry and registrar systems. I'll certainly be suggesting some people back at Nominet when I take a look at this when I get back home, good work.

Just before lunch, we had Fred Baker with who showed his details on various methods of DNS privacy and how QDM minimization works. An interesting approach to DNS privacy, which personally I really support. I think this is a future of DNS privacy at the authoritative layer. After lunch, we continued with Ted and Sarah who gave us an overview of the changes they've made related to hybrid meetings and how things were expanded. Hybrid meetings are clearly the future, so as not everyone can travel an ICANN needs to be inclusive. So these changes are really great and encouraging to see.

Next is time for myself and Donald Rossley to talk about the Customer Standing Committee and the effectiveness review. I'm not going to mention much more about that here other than to mention, again, there's a public comment open. So if you want to comment, I encourage you to do so. We then heard from ICANN engagement team. Yes, it give us an overview of what they do, which was very much more than I thought they did. So that was quite eye opening.

And I thought the details of their Cloudlabs was particularly interesting. After the coffee break, we only heard from Jeff Bedser. Jeff gave us an interesting overview of DNS abuse registry best practices, which is a hot topic at the moment. ICANN is something I'm really interested. So this was engaging content for me at least.

We then moved on to Peter Lowe. Look, Peter talked about how difficult it is to define what DNS abuse actually is and how it can be misinterpreted. This is something many people have wrangled over ICANN over the last few years. Before the last break, Adele from ICANN's engagement team presented KINDNS. This is similar to the internet industry success story manners. And it's something I think will be very good for the DNS industry. But normally, I'll certainly be monitoring this closely and planning to take power. We then had our final brick, which is very welcome because my jet lag is kicking in at this point.

And then after the break, we had two final sessions. Eduardo told us about the tool that technical services have developed for EAI, at least six to TLD zone files and have this list of MX records and does various tests for internalized email addresses on them. It's always good to see how this kind of work is being approached, and the code releases open sources. It's also always good to see.

And then finally, Kim presented on the new version of the root zone management system. As I work for a TLD operator with reasonably large amount of TLDs, we use this very often. And so seeing development in this area is really, really good to see. I'm almost as excited as Erwin is. I'll look forward to using it, and I'm also very excited to hear there might be developments of the technical checks sometimes in as well.

Finally, I'd like to thank Eberhard and the rest of the tech working group team for the excellent content in today's tech day and all of you for attending. I hope we can do something equally as good for ICANN76 next year.

EBERHARD LISSE: That's it. Thank you very much. You're all released to go home or to go and party.

[END OF TRANSCRIPTION]