SINGAPORE - gTLD Technical Operations Lessons
Wednesday, February 11, 2015 – 10:30 to 11:45
ICANN – Singapore, Singapore

UNIDENTIFIED MALE: Wednesday, February 11, 2015. This is the gTLD Technical Operations Lessons, located in the Canning room. It will run from 10:30 to 11:45 AM local time.

FRANCISCO ARIAS: Hello, everyone. We are going to be starting in a few minutes. In the meantime, you may want to get one of these. In the spirit of internationalization, we are going to have presentations in English and Chinese, so you may want to get one of these. I think the English channel is 1, Chinese channel is 4.

Hello, everyone. This is Francisco Arias, Director of Technical Services at ICANN. We are going to start the session. One thing before we start – well, I guess two. First, if you would like to take a seat in the U, you are more than welcome. Second thing, this will be a bilingual session. We are going to have presenters in both English and Chinese. Unless you speak both languages, you may want to get one of these in the entrance. I'll give you a minute to get that.

Okay, let's start. Welcome, again. In this session, we are going to talk about gTLD Technical Operations lessons. We are going to talk about some of the issues we have seen with TLDs and the technical issues. We

are going to have a panel of registry operators that are going to talk about the issues they have seen, how they have been able to fix those.

The intention of this session is to share the experiences that the registries have had so that others can learn. The intention here is not to shame anyone or anything like that. I can tell you that we have seen issues with a number of TLDs from minor to not so minor. I don't think we could say that anyone is free of guilt, if I may say that.

Without further ado, let's start with the agenda here. This is the [short] agenda. We are going to have a short presentation from ICANN on the issues and solutions we have seen so far. This is particularly in reference to the SLA monitoring system; so, issues that we have seen with respect to DNS, DNSSEC, WHOIS, which are the services that we are currently monitoring.

Then we are going to have the panel and a short presentation for each of the panelists, and then we have a section of questions and answers.

With this, I'll turn the microphone to Gustavo Lozano in the Technical Services Team to talk about the issues and solutions.

GUSTAVO LOZANO: Hello. This is Gustavo Lozano with ICANN staff within Technical Services. Let's talk about the SLA monitoring system. ICANN is monitoring what is defined in Specification 10. If you go to the new gTLD Registry Agreement, you will find the [inaudible] specification called Specification 10. This specification defines basically what ICANN needs to monitor, and basically also defines the algorithm of how we should monitor the critical functional services of the registries.

It's important to mention that in this specification, there are two very important things that are defined there. The first one is the emergency threshold. This emergency threshold basically defines when ICANN may invoke the EBERO. It's very important to mention that this is a "may." It's not an automated process. If we decide to send someone to the EBERO, that's going to be a process, or basically, human means are going to make the decision.

The other thing that is defined in Specification 10 are the service level requirements. There is a Service Level Agreement within Specification 10. That is measured on a monthly basis.

How are we doing this monitoring? Basically, we have a software platform. This software platform was developed by Zabbix. Zabbix is an open-source platform. We hired these guys to develop our own monitoring platform. We also have a backup monitoring system that we developed in-house. This is used when we have a maintenance [window] on the primary platform. But normally and usually what we are using is Zabbix.

We have a probe node network. We have our own 40 probe nodes around the world. We try to place these networks in countries in which there are a lot of Internet users. The data centers that we choose to place in these probe node networks have a lot of access to Internet providers, so we have really good connectivity under IPv4 and IPv6.

It's important to mention that, in Specification 10, the algorithm that is there that is defined on how we are doing the monitoring. It's basically designed with the idea that is false positives are mitigated. It's important to mention that we have never detected a false positive. All

of the issues, all of the incidents that we have detected in the system, have been real issues – have been solved, obviously. But it's very important to mention this. We have not had any false positives until this day.

Once we detect an issue – once we detect an incident – we have an escalation algorithm that works. The way this works is the following. If we detect an issue, we send an alert to all the emergency and technical contacts at the thresholds that you can see on this screen. All the emergency contacts – the three of them plus the technical contact, you will receive this alert.

Also at this percentage, we start automated phone calls. These phone calls are initiated by the system, so they are automated. The way it works is the following. We always call the first emergency contact, and if the emergency contact doesn't acknowledge the call, then we call the second emergency contact and so on. We try three times per emergency contact. You can acknowledge the call by pressing zero.

This is very important. If you receive a call and you don't want the escalation to continue, you just need to press zero, and that means that you have acknowledged the call and that you are aware that there is an issue that you need to go and fix.

It's also important to mention that we have implemented what we call code memory. How this works: if the same group of persons already received a call, you are not going to receive another call in the next 30 minutes. So if you are a portfolio registry operator or back-end registry operator and you have issues with several TLDs, you will receive a call, you will acknowledge that call, and you won't receive another call for 30

minutes, because we don't want to start making a lot of calls to you while you are trying to solve the issue.

There is a navigation algorithm implemented into the solution. That means that if you are a portfolio registry and you have issues with several TLDs, you will try to aggregate the alerts into the smallest amount possible. So if you have 100 TLDs and 100 TLDs are failing, you are going to receive just one alert and one phone call.

We also hire a contractor to provide NOC services on a 24/7 basis. The idea of this NOC is to try to get an acknowledgment from a person, a human being. Because normally, when we do these automated phone calls, we have found that some of them are maybe call centers or automated systems. This NOC, the idea is to get an acknowledgment from a human being, and be sure that a human being, that a person, has acknowledged and understands there is an issue.

The NOC can also provide you with information that we have in the system, so if we send you an alert regarding DNS or RDDS, you can call them and ask them if the issue is still present.

Level 2 and level 3 support is provided by ICANN staff. We provide this service on [inaudible] basis.

A lot of registries have asked ICANN if we can provide access to this information. We are working on an API. By around the end of March, you will have access to all of these data points. This is going to be provided through the RRI interface. This is the same interface that you use for data escrow and for the monthly report. We are working on

that, and you will have access to all of the information that is within the system – for your TLD, obviously.

Let me talk about the issues that we have found, how these issues have been solved. The first problem that we have found is that sometimes the registries are not whitelisting the IP address of our probe node network. We start doing the monitoring and we get blocked by the rate limiting or firewall or whatever you have in the middle.

The solution is pretty simple. Please try to whitelist our probe nodes. You can find the IP addresses in the gTLD portal. We change those IP address from time to time when we find a better data center or better connectivity. But we try to keep those static.

Another problem that we have found is that the registries are monitoring their platforms, but they are monitoring from their internal networks. So we call them and we say, "Hey, we found this issue with RDDS or DNS," or whatever, and the response is, "Well, I'm not seeing any issue. Everything is fine." But the problem is that they are monitoring from their own network. There may be firewalls, there may be balancers, and other mailboxes that are in the path of the end-user to your services and you are not monitoring those.

If you are going to monitor, please try to monitor from an external network, or at least try to monitor your external appliances.

We have found a lot of DNS issues. Well, not a lot, but some DNS issues. We have found DNS issues, DNSSEC issues of different kinds. Some TLDs have lost access to their private keys, and we have gone through

emergency routes on updates. We also have some expired signatures, so they forgot sometimes to re-sign the zone.

We also have found issues with the [sign-in] platforms. So, for example, we found an issue with some dynamic updates in a [sign-in] platform that were breaking NSEC records. We have also found issues, for example, with platforms that were generating several signatures per [inaudible] record, and obviously, this was breaking the crypto part of DNSSEC.

The solution is pretty simple: try to test your operational procedures, including your platforms.

We have found issues with the zone nic.<tld>. Most registries delegate this zone to other name servers different from the TLD. We have found that, sometimes, this zone is not working properly, DNSSEC is not working. This zone is important because if you try to access the RDDS service, you go through whois.nic.<tld>. If this zone is failing, the WHOIS service is going to fail.

The solution is also pretty simple: try to monitor nic.<tld> and try to use the same standards that you are using for your TLD.

In the same regard, with whois.nic.<tld>, we have found the same issue. Some registries are not monitoring that the DNS part of the name is working properly. It's the same thing. You need to monitor that this is working right.

We have found IPv6 issues. From the Specification 10 perspective, IPv4 and IPv6 are equal. So when we monitor, if there are issues with IPv6, we consider those issues to be of the same importance as IPv4. This

means that you also need to try to find out connectivity over IPv6 that is within the same Service Level Agreements that you have for your IPv4, for example.

We have found issues with name servers. From our perspective, sometimes it appears that there are a few name servers that appear to be over Unicast. In that case, the solution is try to use Anycast. We have found issues with name servers. That is, instead of sending the actual TLD response or the actual response, they are sending server fails. What we have found is that the zone transfer platform, the platform that the registries are using to transfer the zones between name servers, is not working properly.

If you are going to monitor the DNS, also monitor the infrastructure that is responsible for transferring the zone files.

These are basically the frequent issues that we have found in the system. Now I think that we can start with the panelists.

FRANCISCO ARIAS: Thank you, Gustavo. Just a reminder, for those that were not here in the beginning, this panel is going to be bilingual, English/Chinese. Unless you speak both languages, you may want to get one of these. Channel 1 is English, channel 4 is Chinese.

We have three registries here. We are very thankful that they are willing to share their experiences with us. Let's start first to my left with – and I apologize, I'm probably not going to say your name right – [Wensel Liu]. He started his career in CNNIC as senior developer in 1999, later became the director of the technology department, and since 2009 he is

the Chief Technology Officer at [Keynet] and CDNS. His company is the back-end operator of several [gTLDs]. [Wensel]?

[WENSEL LIU]: [Chinese language]

FRANCISCO ARIAS: Thank you, [Wensel]. Very good points that you raised. I just took note of a couple, here. The emergency contacts, it's a good point to consider for the contracted parties to consider having your back-end provider as perhaps one of the emergency contacts, or have some other mechanism so that you can quickly notify them so that they can take action on the issue.

With regards to [zone] sources of information or places where you can find help in case of issues, there is of course the [Geo TLD] portal where we can help with any questions that you may have regarding the specifications in the registry agreement. Of course, there is also some self-help in, for example, the gTLD Tech mailing list that we have available and where many of the technical operators participate so there is a place to ask questions and get answers from your own colleagues. Thank you.

With this, let's pass the microphone now to Alexander Mayrhofer. He is the head of the research and development at nic.at. He has been the [inaudible] in designing, implementing, and deploying nic.at's [new TLD] registry solution, called "Registry in a Box," which is now operational for nine [new TLDs], most of them Geo TLDs. Alex?

ALEXANDER MAYRHOFER: Thank you, Francisco. Good morning, everyone. First, let me make a note. We discussed yesterday, and we said it was really funny that actually only the new guys on the block are sitting on the panel, while the well-established, big, new gTLD operators are actually listening. That made us wonder why that was the case. But that's a different story. I understand that maybe bigger companies probably have different [paths] in allowing people to speak about issues that they experienced.

Let us look back. The start, for us, into the operational reality of gTLD operations was almost exactly one year ago when the first of our gTLDs started [with the] sunrise. We were one of the first registries to actually have registration open.

When we did all of the onboarding back then, it was a different time maybe because when I did the onboarding for each of the TLDs, the onboarding information request looked different each time we did it. There was a lot of change on both sides, I think of the industry – on ICANN's side as well as on the registry side, as well.

That was also a time when it was not really clear which of the contacts that we provided in the onboarding information were used for what purpose, especially since, to our perception, it was very wise to supply a role contact as an emergency contact. Because people go on vacation. We were told by ICANN, "No, you can't provide a role contact for your emergency contact. We want the name."

That was sort of confusing us, because we thought that, in case of emergency, you really want to address a group of people. I understand

that, as we talked yesterday, there is no clear documentation how each of those emergency contacts is being used. And the system that Gustavo presented before didn't exist back then, so it was actual people calling the registries.

Quite openly, what were the issues that we encountered? The first issue that we encountered is that when we started with the sunrise of our first TLD, the SMD files that we got by the registrars suddenly looked slightly different than what SMT files that we received from the TMCH test interface looked like.

I had received warnings from my programmers, because they said the TMCH did with provide us with exactly four SMT test files. That is crazy. I can't do serious testing. Expect that there will be some problems there.

That was exactly what happened. We had problems with a subset of [SMT] files – parsing them; there was some change to the encoding or something like that and took us one or two hours to fix that problem and deploy into the [inaudible] registry stem, which was scary.

That would have been possible to avoid if the TMCH had provided, for example, a system where you could create your own test SMTs right on their platform – so, sort of a full test cycle with the TMCH.

The other scenario that we encountered is we configured the list of probes that we were using for whitelisting the WHOIS service when we deployed the system. There was no information on whether those lists would change, how frequently it would change, and how we would actually get updates. As I've now learned, it seems to be that this list is also available in the [GDT] portal; however, since we as a back-end

operator don't even have an account for the [GDT] portal, I can't actually access that list.

I understand there are other sources for that as well, but as a suggestion, I would really appreciate if ICANN could consider giving back-end operators at least some kind of read-only account to the [GDD] of the TLDs they are hosting. Because otherwise, as I said, I can't access the information that is put in there.

What happened is actually that, according to the message that we got from ICANN, "Your WHOIS is down." Problem number one, those monitoring messages, they don't really include a lot of debugging information. So to actually get more information, rather than, "We think it's down," involves contacting you guys. That would be an area where improvement would [inaudible].

Then we had to go back and forth to understand what the problem was, because to our perception, the WHOIS server ran fine. We are just returning an error message that says, "WHOIS quota exceeded in case of a certain IP address goes over the query limit." That was exactly what was happening.

The reason, in turn, was that we didn't whitelist a couple of the new probes that were added into the monitoring system. At some point, the monitoring system had more than 50% of the probes going over the query, WHOIS limit and triggered an alert. It was easy to fix as soon as we understood the problem. It was more like understanding that there is no problem in connecting to the WHOIS server. It's rather that the problem is that we are sort of rate-limiting the probes in the wrong way.

The third scenario that we had when we started doing the architecture – almost five years ago, I think – we thought it would be a good idea to combine Anycast with Unicast. The reason for that was because, in case there is substantial problem with Anycast technology as a whole – which means like weird PGP filtering going on – we really wanted to have at least one Unicast server in the delegation set, so in case all of the Anycast goes down for whatever reason, [backing] some router software, we still wanted to have one Unicast instance.

It turned out that, obviously, that Unicast instance was [also] weak point in the name server network, obviously. So we had a situation where, even though 18 of the 20 physical nodes that we used for a TLD to provide name servers were still up, the DNS service was considered down by ICANN due to a structure of the SLA, and due to the fact that a Unicast server is vulnerable. But as long as there is a strong Anycast network running as well, while the Unicast name server is down, it still doesn't impact the service quality.

I also learned that there is going to be an API for the SLA monitoring, which is really, really nice and was something that we asked for from the very beginning. What we are going to do is we are going to use that API to fetch that information from the ICANN SLA monitoring system so that we can detect an error almost as quickly as you can, and remedy it before it even goes into notification. That's something really great.

I'm proud to say that since we did DNSEC for .at since a while, we didn't have any issues with the DNSSEC signing of the zone. One thing that you actually reduced the quality of monitoring that we can do, we are using a sort of end-of-zone domain. The last possible domain name. I think it's

[inaudible] and a lot of dashes. It's contained in the .at zone at the end of the zone.

However, we couldn't use that for the gTLDs, because you were not allowed to put any other name into the zone besides nic.<tld>. So, actually, that sort of name collision avoidance requirements actually made it impossible for us to make sure that the zone was really, completely, and fully generated and loaded onto the name servers. I got advice that we could, of course, do an RSEP for that, but we thought it would really not be [worse]. There is maybe a little bit of conflict between technical requirements and operational, sensible things to do that then collide with requirements from the specification. That is just the most obvious example that came to my mind.

I think that's about it. As a whole, what I can really suggest is make sure you got the right contacts in the GDD portal. Make sure that those phone numbers that are in there actually make sense so that the ICANN folks can reach you if there is a problem. It looks like we are going to get a lot more information from you guys, which is great, because it helps as a whole operational practice. Thank you.

FRANCISCO ARIAS: Thank you, Alex. A couple of points that I would like to mention here quickly. You are right, Alex. This has been a learning exercise, I guess we could say, for both parties, the registries and ICANN. As we go in the way, we learn more things, and now we have more information that is available to you – like the welcome kit, for example. It explains all of the roles of the contacts and where you can get certain information, and so on and so forth.

There are also some miscommunications that now are hopefully resolved. For example, the possibility to have role accounts for the emergency contacts. That's something that is a must if you want to have 24/7 response time.

On the probe nodes, there is also something that we just implemented. There was a request by some registries that were saying, "We would like to be notified when there is a change in the probe nodes." So we'll be doing that. Future changes to the probe nodes, we will send an advance notification with at least 72 hours before the probe nodes enter into production.

As you mentioned, having more of the [inaudible] information, that's also coming. The API for the SLA monitoring system that will allow registries to have access to the information that we are seeing, almost in real time, as we are seeing it. You'll have to wait until the 10%, which will be the first notice that you would get otherwise.

Finally, the other thing, having multiple accounts in the gTLD portal is also something that we have in the road map. I see [Christine here]. It's something that is in the road, but it's coming.

With this, I would like to pass the microphone to David Peall. He has been involved in the ccTLD.sa – that's ccTLD for South Africa – since 2001, as an [administrator] for school.ca and a member of the .sa domain name authority. He joined Domain Name Services in 2010, assisting the conversion to an EPP registry that now services over 400 local registrars with a million domain names.

David Peall is involved in the [inaudible] applications and back-end for Cape Town, [inaudible]. David?

DAVID PEALL: Thank you for the opportunity to share our experiences. Domain Name Services started development on its registry service in 2010, so it's a fairly new EPP implementation. We feel that this gives us an opportunity to look at ourselves from the outside and identify where we can improve.

I'd like to second some of the calls for clearer notifications about what's required when you launch a gTLD, when those reports are due, and the escrow and all that, because it isn't clear from the welcome packs that you get.

In South Africa, it sounds like it's not unique to us – we struggle with IPv6 port and troubleshooting from our transit providers. They don't actually supply IPv6 to their customers, but if you convince them, they will transit your own IPv6. What we found is we've had to build special relationships with key people in those transit providers, since we are able to pick up the phone and phone someone that we know can actually look at the problem. If you drop the support ticket into there, the support machine, it really just goes nowhere.

That came to us trying to monitor our global reach ability. With the Internet being three to four [inaudible] wide but containing tens of thousands of [inaudible], this is a bit of a daunting task.

So we turned our monitoring to point at the ICANN probes, but they seem to be a bit of a moving target. We're looking for suggestions. But I think the API may solve most of that.

Then, to move to the next item, we were first a cc, so we have zone generation servers and we have zone distribution servers, and monitoring around that, making sure the zone is moved from a distribution center to one of our secondaries. When we did the gTLDs, we inserted a blip on the wire, a black box, DNSEC signing engine. When you're not looking, it turns off, and generation is working and your distribution is working but there's no new signatures coming through. We had problems with that.

We've added, obviously, new monitoring, and this is an iterative process to make sure that what we're generating and what we're distributing are the same thing.

A lot of this could be identified a lot quicker if back-end providers had access to an API, and we see that the gaps are closing and these problems will become smaller and smaller. But maybe sharing them up front with other registry operators means that they won't trip over the same things.

The last item I would like to bring up is the very poor response from registrars in terms of TMCH integration. We were very disappointed during our sunrise and [inaudible] that we had customers knocking on our door going, "We'd like to – you guys are TMCH [SMD]. How do we get to you?"

We only had a handful of registrars that even bothered to do TMCH integration; the likes of [inaudible] Energy and MarkMonitor. And they don't really service the public. Perhaps there's an opportunity here in next rounds, or even further applications where registries are prepared to do development and innovate. Registrars are really just there to service the bulk customers. Something out-of-band at the registries could operate to provide claims checks, to provide claims notices, to accept SMDs out of band. It is possibly something we'd like to suggest. That's me.

FRANCISCO ARIAS: Thank you. I think we are starting here in the kind of things that we need to improve. The probe nodes is a common topic here. Hopefully these new processes that we are putting in place will help, so that you will have a [inaudible] notice of the changes we're making. Of course, the API has been a common request across many registries, having the ability to have that information in real time.

Before opening the floor to questions, I would like to pass the microphone to our ICANN CTO, David Conrad.

DAVID CONRAD: Actually, very short – I'd just like to thank all the panelists for being willing to talk about their desires and the issues that they've encountered. I very much appreciate the input, and I think that sort of collaboration with the community will actually enable us to move forward much more efficiently.

As Francisco mentioned, the majority of registries had one issue or another, so there's no one particularly innocent in this room. But we very much appreciate the input that you've been willing to provide. Thank you.

FRANCISCO ARIAS: Thank you, David. Is there any comment that the panelists would like to make, or should we open the floor to questions? Any questions? [Rach]?

[RACH]: Thanks. I wanted to speak to what Alex said earlier about the fact that we were previously getting notices that "Your WHOIS is down," and obviously we think that our WHOIS is up or we wouldn't be carrying on in this fashion. I'm very grateful for the fact that now we're getting more explicit indications that this attempt failed on this day at this time and this is where, and this is how and when we tried it. Because when I send that to my tech folks, they can do something with it. I really appreciate that, and thank you very much for that.

Can we go back briefly to slide number nine? Gustavo, I believe, was talking about the "three tries with call memory." My question is you give three tires to the same phone number, unless that phone number is duplicated for multiple registries; in which case – let me back up. You give three tries to one phone number regardless of whether or not there are multiple registries involved? Is that correct?

GUSTAVO LOZANO: It's a little bit more complex than that. What we have found is some registries list only the registry operator as the emergency contact. Others list, the back-end registry operator is the emergency 3 or emergency 2. Even within the same, let's say, portfolio registry operator or portfolio back-end, we have found different scenarios.

What we tried to do is the following. If the same e-mail addresses – we group all the alerts, and if we found that those alerts showed us the same e-mail addresses for the four of them, then we consider that to be an aggregation point. We will try to do the calls and we send the e-mail alerts and we try to do the direct phone calls for that group. But if you have different scenarios within your TLDs and the e-mail addresses are different, then you may receive several calls; maybe for a group, maybe for a [new] group of TLDs, maybe for a group of TLDs.

Basically, the e-mail address is the information that we are using to do the grouping of the information.

[RACH]: Okay, but we're talking about phone numbers, I think? Not e-mail addresses? So my understanding is that it's the same. If there are duplicative phone numbers over registries, you will still get three tries, is what I'm getting at. Correct? Okay. Perfect. Thank you very much.

ATSUSHI ENDO: Atsushi Endo from JPRS. Please go to the slide number 13, at IP whitelisting. Yes. This says "Full list available in GDD portal," but right now I cannot find it in the GDD portal. Is this my problem or other operators cannot find it, or not?

FRANCISCO ARIAS: If memory serves, you find the list [inaudible]. And also in the welcome kit, there is a section that has a link to find the probe nodes. The welcome kit, or in the [inaudible] in the GDD portal.

ATSUSHI ENDO: This is a related question. On the last year at Registry Road Show gave us the information of the IP addresses, and also in the presentation it's listed and also that – yes, as you mentioned, there is a list of the addresses. We find that it changed. You access "Address not listed in the whitelist," so is it updated? You changed that or not. If in the case it changes, please provide us the information as soon as possible; on the other side of these things, is it okay to block the addresses not listed in the whitelist?

FRANCISCO ARIAS: Yes, of course. The IP addresses that are listed there are the ones that we're using for monitoring. But as we have said before, they have changed from time to time as we find new nodes or change then, because – I don't know, they have some issue in some data center. So that's the process I was referring before that we will be starting with. The next changes that we have, we will give advance notice to the registries before putting the IP addresses in operation, 72 hours ahead. Thanks.

ATSUSHI ENDO: Thank you.

FRANCISCO ARIAS: Joe?

JOE WALDRON: Thanks, Francisco. Joe Waldron from VeriSign. I'm going to make two points, or questions. First, I think I heard from all the panelists some observations that sound very familiar. I'm going to repeat a request that I've made in the past, which is when it comes to the actual implementation of some of these back-end services that we have that ICANN is contracted with, I think it's absolutely important that we have a community review of that statement of work in some form or fashion. Specifically, what the monitoring is around SLA monitoring is important to registry operators.

That is, ultimately, a contractual impact to our ability to continue to operate. It's important that we understand, in a significant amount of detail, how that monitoring is taking place. There's a good community review of that, so that everybody has a good understanding of how that's working, rather than just say, "The information is posted by Zabbix," but there are operational implementations of that that I think are important.

If we want to be on the same page in terms of how we're monitoring and what our expectations are, I think it's important that we have some kind of community review of those systems. I think the Trademark Clearing House is an example of something similar where that kid of engagement would have benefited, and certainly would benefit going forward.

` My second point is, again, to the IP whitelisting. I have some serious concerns with a concept where we have a list that's posted that the registry is supposed to just go accept the IP addresses that are posted on this list and whitelist those. You're talking about access into the registry systems, and this is a security and stability issue for me, that this is a company that is essentially acting as a registrar. We should expect that they follow the same types of procedures, processes that have been long established in terms of how we grant access into our systems, to our registrars. There is a registrar account that they are allocated. We have a process that we use to update our access control list. We should expect that ICANN's registrar that is performing this SLA monitoring would follow those same processes.

FRANCISCO ARIAS: Thank you, Joe. Regarding the documentation, I would say that again a level we have, of course, Specification 10. In regards to the monitoring system, it's an open source option that we are using. I asked Gustavo to post the link to the SBN repository so that anyone that is interested can get the access to the code so you can see, down to the last detail, how the monitoring is implemented.

Regarding the list of the nodes, that's a great point. Right now we have it on STPS. But [if] we would like to have a PGP sign or some other mechanism. I am happy to hear any suggestion on what you would like to see there.

JOE WALDRON: Yeah. I'm happy to share our normal processes that we've had registrars using to request those types of updates. There is just a standard process. They request a new IP address be added and we put that in. That way there is a request, there is an acceptance of the request, and then we implement that, rather than going out and just pulling down a list. I've gone and looked at some of those IP addresses and they belong to companies I've never heard of. To add an IP address that has access into our systems that is questionable I think causes me some concern.

FRANCISCO ARIAS: I will take that as a compliment, the fact that you don't know some of these companies, because that means we are doing a good job in finding diversity on the connectivity that we are using. But thank you. We are open to hearing any suggestions of how to improve this process. Rubens?

RUBENS KUHL: Rubens Kuhl, NIC.BR. To Alexander Mayrhofer's comment of DNSSEC monitoring, I would like to point out that ICANN has already suggested in its registry amendment negotiations the allowance for more records for DNSSEC monitoring. They are [inaudible] to that request.

What I didn't hear, though, in this presentation was about services that are not listed as critical services. So, they are only talking about the Specification 10. It specifies critical services or data escrows. But there are services that are not considered critical, according to an agreement, like CCDS, like zone file access, that ends up being critical. Because zone file access is today what is driving name collision monitoring.

Name collision monitoring is business critical to registries because it ends up defining when a registry can sell domains and can activate domains after a controlled interruption period. We have suffered several issues with the considerate noncritical systems, like zone file, [inaudible] and so forth when IP addresses were not updated or not listed. I would like to suggest that those systems that, because they have a compliance effect and a business effect, would have the same treatment as the [current] critical systems. We recognize that they can have an impact. Even without, they don't call them critical. Thank you.

FRANCISCO ARIAS: Great one, Rubins. [Rach]? And then Chris.

[RACH]: I have another quick question about whitelisting. Would it be possible to get a list of the e-mail addresses and phone numbers that you use so that we can whitelist those as well?

UNIDENTIFIED MALE: Yes. We can provide that list.

[CHRIS]: Thanks, Francisco. I'd like to echo those concerns that Joe has raised, especially around how this monitoring is being implemented. I understand that the software is open source and we can go and download and look at the software, but that's just one part. This thing is now been architected somehow. Somebody has decided frequency of probes and so forth.

You've decided that you're going to have I think it was 40 probes out there. Are they each doing it once every 30 seconds, or are they rotating through? These sort of things would be good to understand. Particularly with some of the smaller TLDs that we're operating; we see that ICANN is actually responsible for most – if not all of the load on the system rather than anything else. Perhaps, once we get a better understanding of that, maybe we can have some discussions about whether those frequencies are actually required for your purposes and so forth, especially on some of the smaller TLDs.

So, general documentation or something that we can have a look at to understand how that works, that would be really helpful.

I then wanted to know if there was a way to opt-out of phone notifications, especially for the earlier stages, the 10%, 20%, or 30%. Especially if something happens at a point in time, what I don't want is the technical people that are trying to resolve the issue being bombarded with phone calls. I want them to resolve the issue. If they don't answer the phone call, apparently they're going to keep getting called until they do. We could probably try and move those phone call notifications to another area of the business, but then it kind of defeats the purpose. I wonder if we could consider that.

Then, with the 72-hour notice of changing the whitelists, I was just wondering, how did we decide on 72 hours? Some registries – depending on what we're updating, some registries are really [quiet]. By the time you go through change control and sign off and so forth for these things, do the checks that you need to do and so forth, it can take longer than 72 hours, particularly if it's posted on a Friday. I don't

necessarily want to have people have to come in on weekends to update things and so forth. I'm just wondering if we could reconsider that a little bit. I think that's all I had on my list to write. Thank you.

FRANCISCO ARIAS: Thank you, [Chris]. So let's see. I got three questions, here. The frequency, that's very easy. That's in [inaudible] time. We query DNS every minute. We query our DNS every five minutes.

JOE WALDRON: Is that every probe every minute? Or are you just reading it across the probes?

FRANCISCO ARIAS: Every probe, not.

JOE WALDRON: So per probe.

FRANCISCO ARIAS: Per probe, yes.

JOE WALDRON: Okay.

FRANCISCO ARIAS: In fact, that's describing the [inaudible]. That's a high-level requirement. The opt-out for phone calls, we currently don't have that feature. If it's

of interest, we can certainly consider implementing. Let's talk more about that.

Regarding the 72-hour, if you think you need more time, I'm happy to change that. So please, propose something. You know my e-mail.

JOE WALDRON: Yes, thank you.

FRANCISCO ARIAS: Go ahead, Alex.

ALEXANDER MAYRHOFER: Another thing that just came to my mind. You are currently not doing the full EPP test that I think you have discussed you would be doing. So it would be actually good if, once you know how the tests are going took like, that you talk to registry operators and tell them what you're actually going to do. So I'm wondering what kind of EPP transactions will those EPP registrars be doing?

We also see the registrars that we have created for that monitoring system are sort of a security risk, [inaudible] before. Because they are sort of idling in our system, the password has never been changed. So we really want to get the heads up before they are going to use [both of them].

FRANCISCO ARIAS: Absolutely. Before we enable EPP, we have that comment from Joe and others. We understand the [inaudible] of that system. Before we enable

that, we'll be talking with you on how is the best way to [inaudible] those things

ALEXANDER MAYRHOFER: Also, short comment to that. It's also interesting because those domain names that the registrar will create will obviously also show up in the ICANN report. Who is going to pay for this?

FRANCISCO ARIAS: There is an easy answer to that. There is a provision in Spec 10 that says there are no fees for those, and I think the plan that it's to have only one domain name, if I remember correctly. So we create it once, and then we just do updates to monitor the system.

UNIDENTIFIED MALE: [inaudible] from Tango Registry Systems. I have three points. Two are same things. First, for the session. Second, that you will provide us with the API for your monitoring system. And my question to this is: did I understand it right that it's the same interface where we put the monthly reports?

FRANCISCO ARIAS: Yes. The API to the monitoring system will be using the same – an extension of the API that you are currently using to send the monthly reports.

UNIDENTIFIED MALE: And then I have a friendly suggestion, that we get a second credentials, because in our case, sending the monthly report is a totally different department than the monitoring department, so perhaps we can get a second credential there, and this is only read-only or something else. This would help to set up different departments.

FRANCISCO ARIAS: Okay, good point. The current plan we have is to – we were trying to minimize the need to have more credentials, so we were planning on using exactly the same users. But I get your point. So we'll look into that. Probably it will not be in the release that we have already planned, because it's already underway, but we will consider for the next release.

Yes, [Chris]?

[CHRIS]: I just picked up on something in the conversation just then, that kind of goes to illustrate the point before about understanding the architecture and how it's going to work. You just mentioned that with EPP, you're thinking of using one domain and just updating that, presumably from all the different probes.

If we go back to the frequency, if there are 40 probes, and if it's set up in such a way that let's say they're doing their monitoring check every minute on the minute, not staggered or anything like that, when you design a registry system, one of the things that you tend to not worry about is multiple commands coming in to update the same single domain simultaneously, because generally that doesn't happen. A registrar typically doesn't do that.

What you're actually going to find – in our registry, anyway – is if you had 40 probes attempting to update the one domain at once, one of those, depending on the timing, is going to succeed and the other 39 are going to fail, saying that that domain is currently being manipulated by another session.

This is why it's important that we understand this architecture, so that we can give you this feedback to do it. I guess that's a purpose of today, but we need some more in-depth detail to let you know. You'd be getting lots of false positives from our system when you turn that on if you used a single domain for EPP. Or you'd have to look for that specific response, but you won't be getting the success response that I assume you'd be looking for. It's a good example of why we need to dive into this a little bit more.

FRANCISCO ARIAS: Like I said, EPP, we're not enabling that yet. We will certainly be talking with you and the registries before we [inaudible] it.

We are four minutes before the hour. Any last comments from the panelists? Any last questions? No? Okay. With this, I will thank the panelists and everyone for attending. Thank you very much.

[END OF TRANSCRIPTION]