STEVE CONTE: We are recording. If we can go ahead and take a pause for recording start. All right. Welcome, everyone. This is day one of the SSR2 ICANN SSR Subteam. It is October 9th. Happy Monday, everybody. I'm Steve Conte. I'm SME for the Office of the CTO and kind of voluntold to be the staff to facilitate today.

Since we're face to face and it doesn't look like we have any remote participants, but we might have some in the observer room – and for the record, since we are recording, I'd like to start by just going around the room and introducing ourselves so we know who's in the room and [who is] starting with that. James sounds like he's wanting to volunteer, so go ahead.

JAMES GANNON: Thanks, Steve. Am I live?

STEVE CONTE: Yes.

JAMES GANNON: Okay. James Gannon, and I need to do my standard legal disclaimers that I'm here as an individual, not representing my company. Any views or opinions that I express are my own and not intended to reflect anything to do with my employer.

DENISE MICHEL: Denise Michel, I'm a co-Chair of the SSR2 full team.

YVETTE GUIGNEAUX: Hi. Yvette Guigneaux, ICANN staff.

NORM RITCHIE: Norm Ritchie.

BOBAN KRSIC: Boban Krsic, member of SSR2, rapporteur of the subgroup ICANN SSR.

ZARKO KECIC: Zarko Kecic, member of SSR2.

NOORUL AMEEN: Noorul Ameen, scientist from CERT India.

NEGAR FARZINNIA: Hello, everyone. Negar Farzinnia, ICANN staff.

STEVE CONTE: And just noting that we have Maguy Serad, not next to a microphone, just observing the opening proceedings this morning as well in the room.

So, welcome. Does everyone have the Wi-Fi password, and everyone in? Because I'm going to change the slides. Okay, so while we're doing that, I'll do some other housekeeping here.

For those in need of biological breaks, if you go straight down the hallway behind the screen, and the first left will get you to the restrooms. If you're looking for caffeine or [inaudible] and waters are behind you. If those are not, we have a café just on the other side with cold drinks, so we can get you guys over there.

We'll go through the agenda in a second. Just to set the stage a little bit from ICANN perspective, this will be somewhat interesting for us. We have I believe only two people who have signed NDA, but I don't want that to stop the discussions this week.

So we are very interested in having these discussions and working through the work plan, and I ask for patience as when questions come up and we have the speakers with us, any silences that might be perceived as awkward silences are actually going to be silences that are trying to frame an answer that would fall within acceptable outside-of-NDA type answers.

If it's not an acceptable answer, then we can start working on asking the question in a different manner or trying to determine where the sticking point is. We want to be open, we want to be contributory to this meeting, but we need to fall within the rules of our own staff employment too in protection of ICANN. So I ask for your patience and your willingness to work with us on that. Zarko, you had a question around that.

ZARKO KECIC: It's not a question, just a statement that I'll sign NDA as well. I printed it out, but my legal department took that and kept that, not to observe but to copy some good stuff from there. So I signed it.

UNIDENTIFIED FEMALE: [Anyone else?]

STEVE CONTE: Denise?

UNIDENTIFIED MALE: I signed mine, just haven't handed it in yet. Let me do that right now.

DENISE MICHEL: My company has not yet completed its review of the NDA, so I don't know whether or not I will sign it. But I'd like to suggest a couple of perhaps working methodologies for this. When issues arise for which staff feels that an NDA is required or there's confidential information, I think in my mind, there r a couple of classes of information here.

There is information that can be discussed openly with the team and the public. That's Class 1, and Class 2 is information that the team would discuss with Review Team members but does not want to discuss publicly or to appear on a public list. That's Class 2.

Then Class 3 – and I'm just making up these classes – are true NDA covered information that you do not want to disclose to any individual outside of ICANN without a signed NDA. Would you agree with those three kind of buckets of information?

STEVE CONTE: In general – and we'll find out in practice if that works – I think there's not a Class 2, and frankly speaking – and I'm not IT, so I'm proxying them right now – we're not going to show areas or discuss specific areas where there could be a compromise or there could be potential for compromise, because that exposes more risk to ICANN as an organization.

So I ask that the Review Team – I don't think we're going to get to a point where we're going to, say, list the passwords for our routers and things like that, and that's just stuff that even under NDA, we wouldn't disclose. It's stuff that I don't have. So there'll be a certain level that even under NDA, we won't disclose. And if we get into that point, then we'll also be asking probably – if it gets to that granular level, we'll probably be asking how it relates to the work that's being done with this Review Team.

DENISE MICHEL: Thanks. I appreciate that. If I may follow up, of course the first Security and Stability Review Team whose work we are reviewing and assessing did not use an NDA to complete their work. Every Review Team of course travels its own path, but I would ask that if this group goes into the territory that's addressing Class 2 or Class 3 areas, I think it's

important that staff be clear about that, and that we get a follow-up note saying, "We were asked for this. Here is the restriction that we placed on this information, and here are the actions that need…" So it's just really clear for everyone whether an NDA is needed or whether we need to discuss this outside of microphones and public e-mail lists. I think it's important to make sure that we keep track and document where we do run into the situations where there is less than full and free exchange of information. Would you agree with that?

STEVE CONTE: From a staff level, I would agree with that. I'm looking at the whole team to see if everyone is nodding in an affirmative or – I see no fist shaking, so I think we're okay on that. Eric, welcome. I heard you went for a scenic drive this morning.

ERIC OSTERWEIL: [inaudible]

UNIDENTIFIED FEMALE: [inaudible]

ERIC OSTERWEIL: I'm new here.

STEVE CONTE: Final bit from me on this too is just – I'm hoping that the NDA is a nonissue, and I'm hoping that we can have as many open discussions as

possible and get the answers – the intention of the answers that you guys are looking for, without having to fall under NDA.

The fact that more of you are signing is fantastic. It doesn't mean that there's a magical cart that's going to come down the aisle with all of our sensitive information, it just means that we're going to have a deeper dialog about a specific topic, and there's no way that I can pre-flag that so we just have to go as we move forward.

With that, I think unless there are any questions from the Review Team to staff at this moment, I'm going to pass it on to Boban. As co-Chairs, he's indicating – do you guys have any messaging that you'd like to give to the Review Team prior to begin? I'll pause, and if not, then Boban.

BOBAN KRSIC: Hi. Yes, first of all, I want to thank you. Thank you Review Team members to make this happen, to be here today, to assist in this topic. We have a lot of staff here. I would also like to say thank you to ICANN staff and to organize it, and of course the people who are here at the time with the agenda.

From my perspective, it isn't a classic audit [inaudible] so we should try to go into a dialog, and not only monolog. We don't want to find any issues that are here and say, "Okay, you have problems," or what else. So we should address the topics there. We have four or five topics here.

We have something about business continuity, something about incident response and security management, something about new gTLD allocation and transition processes, and we will take a look at the

process side, on the procedures, how is it documented and how you live it, and how it works.

So that's the base, but we don't get on the detail level yet to [inaudible] and we have possible recommendations here, and are they here [for fit in] assistance. Yes, that's not our scope here today. So it should be open, and we have a lot of stuff here.

We have a lot of different topics, and we have only two days, but I think these two days give us the opportunity to move or to understand it better for what is ICANN here responsible on this topic or the SSR of unique identifiers, and how can we help to make this better. So it's mainly just understanding to get the understanding, and then to work on it. That's it from my side. Thank you, Norm, also to participate here. Eric, Denise, great, and also the other team members.

With this and with a look on the agenda and the first topic, we'll start. I don't have my glasses so I don't see this, and my MacBook doesn't work, so yes, [inaudible] change.

UNIDENTIFIED MALE: [inaudible]

BOBAN KRSIC: Yes, sure.

JAMES GANNON: Just very briefly, I want to build on something that Boban said. And I think there's been some – I'll probably say miscommunication on our part about the "audit" piece. So I think a lot of this may have come from me and Boban wanting to put structure around – I think somebody called it a fact-finding mission, and I think that's a much better way to phrase it.

We looked at two industry standards which are security management and business continuity management as a way to put structure around the topics that we feel are important to address, but I do want to make it really clear that the intent here is not to audit ICANN. Under no circumstances is that the intent.

What we want to do is use those scriptures as a way to guide us so that we're not kind of free flowing it, and it puts it then more a way to say, "Okay, here's why we're looking at that based on that audit structure," rather than, "We're here to check your password compliance policy" or [inaudible] you. That's not the intent. The intent is to use a structure to be able to guide us in conversations with staff, because I think with only two days, that's the only real practical way to do it.

STEVE CONTE: Thanks, James. I appreciate the clarification. I appreciate the comments from Boban too. I think the confusion – I could go both ways. I'm not willing to agree, I'm willing to stand with you on the confusion. I think if we're careful with our lexicon this week, we can go a long way.

Those of us who joined the calls and have been on the calls every single time, we understand what the intent and the intention of the mission is.

We are inviting members of staff who haven't been involved on a regular basis with this, and I think if we collectively – staff included – are careful with our lexicon, I think we can get past any roadblocks that we might have built ourselves.

I ask also to the Review Team as we go through this, and I think this was in your message too, is that – especially because one of our first sessions here is subteam face-to-face discussion. So just to help set the stage, and then I'll really stay away from the microphone and let you guys discuss, but I ask as you guys reflect and prepare for the two days that we look at – when you look at your work plan, you look at how the questions you're asking, how it's going to affect the SSR of unique identifiers.

And in this case we're talking about ICANN as an organization, and there's certainly impact on that, and I think there are some valid questions there. But the questions in my opinion – only my opinion, the questions that should be asked, the discussions that should be going on is, do the line of questioning or the direction of questioning relate back to the unique identifiers and ICANN's role in managing those? And go from there.

Boban, can I just pass it on to Eric? And Eric, I'm going to ask that you use a mic. We have nobody joining us remotely, but we are recording this.

ERIC OSTERWEIL: Yes. Thanks, Steve. I just [inaudible] the record need my 405 comments. But for the record, I object to taking the 405 in the mornings. That was a

previous comment that [was online]. No, so just to pile on to what you're saying, Steve, I think our intent here is to sort of assure our review that there isn't really anything that we need to be looking at deeply. I don't think that any of us presume that there's a purview to go digging deep. I think what we want to do is ensure that the identifier space is a separate piece or that there isn't sort of a business continuity issue that really is something worth talking about.

We want to basically put those fears to rest, but we have to look to see. I think that's kind of – [not to] put words in the rapporteur's mouth or anything, but I think that was a lot of our perspective.

STEVE CONTE: I would agree with that, and just would say that we are not at perfection. ICANN as an organization. I think no organization can claim perfection. I think if there are areas identified that we can work on, I like to believe that we're all here because we want a better ecosystem on the Internet, and so we're all working towards the greater good, we're just coming at it from slightly different directions. So with that good faith statement, that's how at least my intention on operating in the next few days will be with this team.

BOBAN KRSIC: Yes. Thank you all. Do we have comments from the team on this? Yes.

JAMES GANNON: Just very briefly, I think to again put the same expectations to the staff members who maybe haven't been following our work, I think the first

questions we should be asking – particularly those who are not familiar with what we are here to do – is within your domain of knowledge or domain of influence, what is being used to coordinate the unique identifiers?

So for example, I notice we have a VP of Risk Management coming in. Okay, so there is going to be a split there in his roles and responsibilities with regards to, for example, business continuity for the PTI functions which are provided under contract. That's something that's definitely within our domain of what we need to look at, but whether somebody from an accountant from finance can access from home in the case of the office being shut down is not something that we need to necessarily be looking at. So I think we need to give them that kind of sphere to work in.

STEVE CONTE: And even just [inaudible] example too, if we are talking about risk management and stuff like that, maybe this is a discussion point, but my sense is that – these are my words, so don't go eat into them – I don't feel the relevance would be on if you're talking about the risk management plan on who we contact, who we contracted with outside the organization in case of an event of some kind, data or physical. It's not relevant.

The fact that we do contact, that we do have a process in place is more relevant than the whos. My opinion, so if you guys disagree with that, please let's have a discussion about that. But I think that's going to be a

point where I think if we can agree with the level and intention of answers, I think that's going to help [inaudible] the next few days out.

BOBAN KRSIC: Related to this, my perspective or my opinion on this is we should start all the way at the process description [inaudible] with the procedure, how is your risk management process in general, and are the mandatory steps – and the mandatory steps are described in industry standards – are they there or not?

We don't want to look at the output of your risk assessment. So your risks are your risks. I can only say, "Okay, is the process there? Is the process documented and are you working like the process as described?" And I think you have the expertise to say, "Okay, that's an issue for us," or "Not an issue for us."

So for me, it's only the first part relevant. Is the process over there as he described? And it's a full process like the industry standards, it's defined.

STEVE CONTE: All right. Well, we're into discussion of the subteam at face-to-face discussion, so I'm actually going to sit back. And if you guys want to talk – however you guys want to talk about it, I've got one physical copy of the work plan here. There is a link on the agenda that Jennifer sent late last week, and I think Yvette is going to bring it up as well.

I'm going to pause and see if there's any strategic way or tactical, or whatever you guys want to do on how to approach this. We will then go

into a break at whatever time that was. 10:45, a break there, and then at that point, we'll collect James and Xavier and we'll have the first session with staff. So at this point I'm going to sit back unless there are questions to staff or anything that you need me for as a [gopher]. Just let me know.

JAMES GANNON: My idea is that – people feel free to disagree with me – we use these questions as a way to guide a discussion. So at the end of the day, the people who are coming in are the SMEs in their area, and they know how an organization is operating within their area – or I hope they do – and by using these questions, we give them a structure to give us, at a high enough level that we're comfortable but at a deep enough level that we're getting a feel for the actual execution of these areas. That's going to be the balance that we'll have to find with each speaker. That's fine. But we use them as a way for them to guide us through your processes as ICANN organization. Because at the end of the day, ICANN does risk management, to use that as an example again. That's a huge, broad statement. We want to understand – within the remit that we have – what direction, what strategy and what approach is ICANN taking to that goal of being a risk managed organization.

And by having these questions from an industry standard which most of these speakers are going to be aware of, ICANN may not be following it formally, but these guys are going to be experts and professionals in their own areas and they'll be aware of it, that it then allows the speaker to give us a sense of – let's take the very first one. Does ICANN use formal [ISMS]? Is there a security framework? How are you

executing against the security framework? What items within that are potentially missing? What are the struggle points?

One of the things I'll probably be asking most of the speakers is, "What do you feel is not working? What do you feel is – not necessarily a gap that's there, but what is something that needs to be built on over the next three to five years, for example?" Because that's where I feel the Review Team has the ability to come in and enable ICANN staff through a recommendation to be able to go further in their own debarments and goals, because there may be things that some of these guys want to be doing that they feel more attention should be put on in the next few years, and that's where the Review Team may choose to make a recommendation that, yes, that is something that ICANN should be focusing on. And by having those discussions and conversations in the bounds of these sets of questions, I hope that we surface where things are going well, where things are potentially struggling. And then hopefully that will then turn into recommendations coming out of this one.

BOBAN KRSIC: Thanks, James. What do you think about when we start with an explanation of the process when the first comes? It's a good way. Let them explain how it is implemented, how they deal with it, and then we go into conversation. Because I don't want to start with questions. So put them that they talk and give them the chance to explain it. And then we will try to address open issues from our question list through the dialog, would be an appropriate way.

NOORUL AMEEN: I support with your review point. Let's understand first, then we will have a discussion.

BOBAN KRSIC: So before we start later on, I've updated Trello with your interest you signalized, where you have your expertise, and put [there] some names on the relevant topics. Only to be sure at the end that we have responsibilities related to the topics and that we have two or three people who are working on this and who will write recommendations down and put them into the final document.

My idea is – sorry. [inaudible]. The idea is to take a look at the Trello board and say if you're comfortable with it or not, and then that we decided, "Okay, who's the responsible team member for this topic?" James.

JAMES GANNON: Maybe it might be an idea to pull up the Trello board and we spend 10 minutes going through that. Yvette, would you be able to pull up the Trello? Or Boban, do you have it? Okay. Or if you give me presenter, I can put it up. One moment, please. Technical difficulties. Okay, Yvette, can you give it to Norm? Because mine is done.

DENISE MICHEL: I can fill some air time as we work on the Adobe. I haven't gotten Trello added on my business laptop yet. I'm still hopeful, but I just wanted to

note that I have some specific suggestions for additions to flesh out what I think we should review, and would really like to get your input on to flesh out some other aspects of this plan that don't relate to the information security management systems. And I think one of the areas is – and again, this flows from some of the recommendations that we have to look at from SSR1 as well as other topics that we've identified as of interest and within scope.

That is I have some contributions on ICANN Compliance for the team to consider. I had some suggestions on looking at the scope of ICANN's SSR responsibilities. I think we've already noted business continuity management. I think it would be helpful to flesh out within that we should be looking at. I think those are the key areas, and I have some sort of specific questions that I can send to the group that flesh out those three areas in particular a little bit more.

And then I'll get some help to turn those into Trello tasks I think is what you would prefer to have shown there as well as we'll add them in a different color to the list so people can also consider them that way. And again, these are suggestions for the subteam's consideration. Thanks.

JAMES GANNON: I'll fill some airtime as well. With regards to your suggestion, Denise, are you suggesting – one of the tasks we have is looking at WHOIS ARS, [inaudible], EPPD, the UDRP, registrar data escrow, so the ICANN systems involved in managing those are on our list already. Is this on top of those, or are you looking at the actual implementation of contractual

compliance? Because that I think is a more complex discussion that I think we'd need to have.

DENISE MICHEL: It's on top of those, but the compliance portion I think needs to be fleshed out and provide more detail. And I think the group needs to make some initial decisions about what within compliance and the intersection of compliance and SSR2 we want to focus our efforts.

NORM RITCHIE: Have we discussed how the ccTLDs fit into this? It's kind of an odd duck because it obviously is not under contractual part of ICANN, but if you look at the root DNS, obviously the ccTLDs are there, and they have a strong affinity to IANA, whatever it's called now. So I'm just wondering if we've already discussed how we're going to deal with that, or just ignore it. That's probably something we should decide at some point.

STEVE CONTE: I don't think the team has spoken specifically about the CCs, but there was discussion at the very beginning of the different boundaries of "ICANN." There's ICANN as an organization, there's ICANN as a community, and the different aspects – especially when we're looking at DNS, the different aspects, such different pieces of that.

And I'll leave it to the team to talk about the specifics of your question, but one of the things that can come out of that is if a recommendation comes out of that, just the knowledge from the team that the recommendation is for us to suggest or for us to urge the community,

whatever community that is, CC or other, to do something that ICANN might not have the authority to make a specific recommendation change or an implementation. So as you look beyond ICANN the organization and into the community, I just ask you to reflect on those pieces.

NORM RITCHIE: Yes. That's why it's kind of the odd duck, because obviously, it's not part of ICANN the organization because everything is contract within ICANN, but the root zone is still there. So it's a fuzzy area, and there have been attempts to formalize that in the past, such as letters of affirmation I think they're called, which had mixed success. But I just wonder how this group – we have to address that at some point, can't just ignore it.

DENISE MICHEL: That's a good point, Norm. In terms of looking more holistically at the context of the DNS and several of these parts, I think ccTLDs definitely come into play, but we haven't articulated specific review actions in this area, and really setting aside the ccTLDs contractual relationship I think with ICANN, there's certainly – as you know – other connections. But it's really just not an area that we've discussed much at all.

NORM RITCHIE: Okay. We also have the opportunity to ask them. We can also do that.

BOBAN KRSIC: Okay. Thank you. Back to the suggestion, Denise, on Trello board we have only a high-level view on the tasks. So as you can see, we have eight tasks here, and every task is only represented in [topic here.] I'm not sure if we should put it here inside, or I don't know how to deal with it with your suggestion.

One of the other things, we can write it down and track it, but it's only a high-level view, and – yes, okay, then

ZARKO KECIC: Yes, this is high level view, but we have some tasks or questions or whatever we called in our work plan, so we can adjust to that and [inaudible]

So yes. What I would like to see or hear is that we should assign ether the entire group to work, or if we don't have enough time to have a team of people working on a couple of issues which are interconnected. For example, business continuity and the risk management are interconnected, so we should have the same group of people working on both of them. We'll have to think about that to see the plans and stuff.

BOBAN KRSIC: Okay. You propose to write your name here? And James, Kerry-Ann. I'll give you a lead in this topic. And now we are fine with what I'm doing? Okay.

So as you can see, we have business continuity, and we have here risk management, and the team members are the same. So you, me, Kerry-

Ann. And you are right, risk management is the first part of the business continuity management, so I want to link them to have the same team members inside. That's what I did.

When we take a look at the security incident response processes, there are – Ameen and Norm, are you comfortable with it?

UNIDENTIFIED MALE: [inaudible]

BOBAN KRSIC: Okay, that's wrong. Risk, risk, security. And these are questions – yes, that's risk management. Let's just copy paste [this]. Zarko, why is the password in here? That's better. A bunch of these are questions that we have here. Incident response time, processes and resources – are they documented, and so on. Are you fine with this? Yeah? Okay. Norm?

NORM RITCHIE: That's a broad question, Steve. Is ICANN ISO 27001 certified?

STEVE CONTE: To my knowledge, we are not. And that's part of what the discussions that will take place, is that when this first started – and James, Boban, maybe you guys can jump in and facilitate or contribute to this – we started a discussion in Madrid in May, and we started talking specifically about ISO 27001, and through discussions that Boban, Zarko, James and I and others have had, we've acknowledged – or I'm hoping, maybe not

– that there are other processes in place, especially within IT and other things, but they're not 27001, and that's why they came down from an ISO-level audit to more of a work plan on understanding what processes are involved.

NORM RITCHIE: Oh, yes. I'm not implying that it should or should not be, it just might really accelerate things if you already have this certification.

BOBAN KRSIC: Yes, it would be easier for us. Just comply to that, and that's it. James, ICANN's processes. New gTLD delegation processes. We have on the agenda I think today – yes, this afternoon two topics to address this item here. Do you have any suggestions here, or are you comfortable with it?

JAMES GANNON: No, I can leave that discussion. Yes.

BOBAN KRSIC: Also, with your leap year as I labeled it, as you can see, we have four team members – so Denise, Norm, myself and you. And I'll give you the lead on this topic. And first one.

ERIC OSTERWEIL: Sorry, just a clarification. So when we're dividing into team members, is the expectation that that just – the responsible parties for drafting, etc.? It's not like an inclusionary [inaudible] make sure.

DENISE MICHEL: Information security management, and can you just be a little more specific about what you're putting in that bucket? My answer would be yes, and understanding it's a really broad area, I think we need to do a little bit of work to be more specific. Just by virtue of being required to look at the SSR1 review and their recommendations and what was implemented, and what the impact that then gets us into infosec. And so I think it would be worthwhile to spend a little bit of time recommending how we draw that circle.

JAMES GANNON: Yes. In the context of what we've discussed before, an ISMS, an information security management system – such as basically every area of the business I hope – and if you look at the speakers in the areas that we're going to be looking at, I don't think we need a dedicated session for infosec topics because – well, I certainly will be asking relevant questions during each of those speakers' areas on that topic, because it kind of goes across all of them rather than being a session on its own, in my opinion. Or that was my expectation.

ERIC OSTERWEIL: To sort of maybe put a fine point on some of the stuff James is saying, I think one of the things we can do if we want to keep this here or have it

as a bucket or a session or whatever is we could use it potentially as a way to discharge our concerns. So infosec is really broad, but one of the things that we talk about is like network segmentation.

So for example, if a set of systems that fall under the purview of an infosec posture are segmented from things that relate to the business continuity of ICANN to PTI or something like that, we can basically discharge our interest in them by saying that there are controls in place, there are audit frameworks in place or something.

This could be sort of the catch-all saying, like we no longer have to look at payroll. We're good with that now. But that's just one way we could use it, because technically, if there was like no network segmentation and everything was in sort of the shared broadcast domain, we would potentially have concerns about a lot of things just because they can touch each other. This could be a bucket for us to start excising those if we want.

UNIDENTIFIED MALE: Just to throw in the final unspoken part as it relates to the management of the unique identifiers of the Internet. Thank you.

ERIC OSTERWEIL: Yes. Eric agrees with that.

UNIDENTIFIED MALE: So the system that I put my request for time loss in doesn't impact the unique identifiers as a clarity statement on that. I hope, yes. If I'm the weak point, then we've got some serious discussions to take place.

JAMES GANNON: Yes, and I think that's the key of how we frame our questioning, is that we're not going to be asking, "Is your ERP system segmented?" It will be, "Is the systems that touch the DNS and the systems that are responsible for the coordination of the DNS segmented?" Once we have that answered, to be honest we – well, some of us do, but – don't really have to care about the rest, business-side systems it's called, but the DNS-side systems, yes. They're in scope for us to ask questions about. But once those questions are answered, to me anyway, that's where our line stops.

STEVE CONTE: While you're on that screen – while you were just on that screen – we have a segment for DAARs this afternoon, domain abuse. If you could just give me – and John Crain is in the air and Dave Piscitello I think is already travelling as well, so [inaudible] for that. I've had a conversation with John last week, and happy to discuss it, but wanted to understand more of what you're looking for in this respect, because we did do that demo in May in Madrid, and I was just curious on what angle sounds [nefarious.] But it's not what angle you're looking for in the question and discussions on DAAR on that. And Norm, I'm sorry, you weren't with us in May, but if you do want to catch up, there's a recording on the

domain abuse reportings that we've put together in the archives in the subteam.

NORM RITCHIE: Thanks. I'm familiar with the –

STEVE CONTE: You are. Okay. Good.

BOBAN KRSIC: Perfect, then we can skip it.

NORM RITCHIE: No.

BOBAN KRSIC: Not really. You want to see it.

NORM RITCHIE: [inaudible]

STEVE CONTE: Yes, but we won't be doing another demo today.

BOBAN KRSIC: [inaudible]

STEVE CONTE: But there is an FAQ that we posted, and we can talk about the strategic implications of how the organization or how the community might want to use the domain abuse tool too. So I'm happy to have that, I just wanted to understand in what respect – since you had it identified – you wanted to look at it from.

JAMES GANNON: So which topic is that related to?

UNIDENTIFIED MALE: [inaudible]

JAMES GANNON: Okay. Personally, I'm confused about where that fits in our plan, but I didn't draft the agenda.

STEVE CONTE: In the work plan on Trello, it was grouped up with some of the New gTLD Programs. EBERO, stuff like that.

UNIDENTIFIED MALE: [inaudible]

STEVE CONTE: It was there. There it is. It's there.

UNIDENTIFIED MALE: [inaudible]

STEVE CONTE: Either way, if you guys think of questions or want to look at it from a strategic perspective on how it impacts ICANN's use, I'm happy to discuss and have that discussion. I just wanted to make sure we knew what we were doing with that.

ERIC OSTERWEIL: Strictly sort of backtracking, I'm not sure this is where it came from, but it could be that that was brought up in regards to when there is domain abuse, how does that factor into ICANN, and then related to any potential action taken in the root or something like that. I don't know if part of the thinking was if this becomes part of the processes, then knowing which systems are involved is important for business continuity. I don't know if that's what it was, but [we could do that.]

STEVE CONTE: Okay, so maybe we keep it up on the agenda for now, and if there are questions or line of thinking that comes up, happy to have a conversation, happy to answer as much as I can, happy to record those that I can't. And if there are no questions, then that's the last item for the day and you all get to go back early.

BOBAN KRSIC: Okay. Back to this issue, we have only two team members on this one. Is anyone else interested here to run through or to assist in this one? We have only Ameen and Norm here. And I give the lead to Ameen, because we don't have any other issue open or [if you] said, "Okay, I'm not comfortable with it."

Yes. I think you write it down in Madrid, so you're the perfect candidate. Not only –

UNIDENTIFIED FEMALE: [inaudible]

BOBAN KRSIC: Perfect, yes. [inaudible] Denise. Ameen, what do you think about it? Do you want to have the lead on this topic, or should we –

NOORUL AMEEN: [It's better if] Denise can lead the process.

DENISE MICHEL: On this one? Oh, I think –

UNIDENTIFIED MALE: [inaudible]

DENISE MICHEL: I think Norm would be a great lead for this group. Are you asking for my suggestion?

BOBAN KRSIC: So Norm, it's up to you.

NORM RITCHIE: Yes, sure. I think this is one of the areas that people probably get nervous about [inaudible] scope of it just because of the title. So yes, I'm happy to take this, but I just want to make sure we have this contained and focused properly.

ZARKO KECIC: And you're welcome to change the title.

NORM RITCHIE: Yes, I think that's what gets people nervous.

BOBAN KRSIC: Then let's do it. What do you think how we can give it a better name? No, it's not only legal.

JAMES GANNON: [System] supporting abuse monitoring, something like that?

BOBAN KRSIC: It's not only abuse monitoring [inaudible]

ERIC OSTERWEIL: This is a very generic set of issues that don't look like they have a lot in common, yet I think they're important. I hate to say it, but until we start to draft something, that probably is the right title, because that's the only way I can see to tie them together. Just put like "working title" at the end in parentheses maybe, or something.

NORM RITCHIE: Yes, I think if we just threw SSR in there somewhere in the title, it would alleviate some fears.

DENISE MICHEL: Sure. This is an area obviously that I work in, so I understand all of these elements. I think there are other people around the people who do as well. I can certainly take the lead on this area, but invite anyone else who's interested to do so as well. A very long yes, isn't it? Sorry about that. I'll have more coffee and be more succinct.

ZARKO KECIC: I have one question regarding this topic, because while doing that breakdown of questions and subtopics, you run into the problem – because most of the reports are talking about WHOIS accuracy and stuff like that, but there are not other issues that may be more interesting for security and stability of the gTLD and their operations.

So I don't know how to find out – actually I look at some compliance report from 2016, I think, and there are still some question marks. So Steve, do you have any information after that period of time, and do we have something more than WHOIS accuracy and some high level compliance from the report?

STEVE CONTE: I don't have that information. That's something we can be asking the staff as they come in. I don't have a good answer for that, so that's ongoing conversation this week on that.

ZARKO KECIC: Okay. Thanks.

JAMES GANNON: I'm just going to put a pin in this topic also that we need to be very careful that we look at the implementation of the policies rather than the policies themselves, because that is out of scope for us 100%. So just as we're walking that line, let's be careful.

DENISE MICHEL: And then I think we need to be really mindful of the community's policy development responsibilities of course in this area and their ongoing work, but we also need to be mindful that a brand new community Review Team solely focused on WHOIS and RDS has just started, and so we want to make sure that we aren't duplicating. And I think we'll need to note things that – and probably have some correspondence with that

team, noting that there is a WHOIS component to the first SSR review, and I think we should discuss how we want to flag things for them and sort of saying, "Here, we assume you'll be taking care of this space."

But I think there still are some important connections to our work that we need to address.

JAMES GANNON: And also on the RPMs, so [UDRP] and everything else, there is also a PDP ongoing for those as well.

STEVE CONTE: In that light too, I don't know how many of you follow the CCT Review Team, but they're coming up with a new set of draft recommendations. In those draft recommendations, there are some areas that if I'm remembering the draft recommendations correctly are suggesting that SSR2 take a look at, or some future SSR Review Team too.

So I think it's perfectly natural to do handoffs or suggest that it's the work of other Review Teams on that.

DENISE MICHEL: I commit to making sure that we are plugged in to all of the relevant activities within ICANN, be it policy or Review Team work, and I'll make sure to flag that and raise that to make sure that this group and the rest of the team understands where all those connections and collaboration opportunities are.

JAMES GANNON: Okay. So abuse is a pretty broad topic. And again, I'm late to the game with this group. I'm assuming there was discussion on how abuse and abuse mitigation fits into the SSR. So, yes, no, kind of?

UNIDENTIFIED MALE: [inaudible]

JAMES GANNON: Okay. I know the CCT Review Team has actually done an abuse study. They've put quite a bit of work in this area, actually.

UNIDENTIFIED MALE: And to that – I'm sorry, I was gone for a couple of weeks in the last month. Did that presentation ever happen in SSR2? I know that Brian Aitchison was trying to get some time on the plate. Okay, so we'll capture that, that the DNS abuse study that was done for CCT, we'd like to get that on the plate for presentation to the plenary of SSR2, not just the subteam.

DENISE MICHEL: So the SSR2 members were invited to participate in a global Adobe Connect call briefing on the report. That's I think something we need to circle back on to see if the majority of interested Review Team members were able to participate in that. And if not, how we want to handle an additional briefing for the Review Team.

That, and then we also have an outstanding request for a briefing on the proposed registry security framework that was posted for public comment. That's still outstanding in our briefing request list.

ZARKO KECIC: Okay. Can we close this topic? Denise, you're leading that. Thanks. What is the last one? Is it this one?

JAMES GANNON: May I suggest that we break early, if we've gotten through everything that we need to? Or [inaudible] Go ahead, Norm.

STEVE CONTE: I am too, but I also have an alternate plan. We could either take the time this afternoon, or Norm had brought up the DNS abuse piece of it. We can have a minor conversation about it, because there are some thoughts from our perspective, from SSR Team within OCTO about DNS abuse and mitigation and things like that, that we can at least open up a conversation, have some thought and then maybe circle around this afternoon to continue that if you guys want, or we can break early. The challenge on breaking early is I can't guarantee that our speakers will be able to adjust to the slot of if we move. I'll leave it up to the team.

ERIC OSTERWEIL: One thing about taking us to DNS abuse is that it's possible other people who aren't on the subteam in the general team would have been

interested in that discussion, so we risk having to just rehash it or backtrack if we come to any kind of [inaudible] on things.

STEVE CONTE: Well, your rapporteur has already broken early.

UNIDENTIFIED MALE: [inaudible]

STEVE CONTE: We can go, we'll go and ask if the staff can adjust. If not, we can break and then you guys can do day job stuff for a couple of minutes or whatever until we can get them in. So officially – what time is it now? It's 10:05. It looks like we're doing 15 minutes. Am I reading this right? Yes. 15 minutes, so if we call it 10-20 to be back, and then we can work out from there whether or not we can get the speakers in or if they should continue on and stuff. Negar?

NEGAR FARZINNIA: I'll definitely look into it. I believe James and Xavier had a hard timeline to start at 11:00, but I will check with them to see if that has changed today or not.

JAMES GANNON: [inaudible]

STEVE CONTE: Yes, we can use that as a filler at any time. I'm not talking about… yeah. So let's come back to 10:20 and we can make some choices from that. All right. Boban, we're on a break now for 15 –

BOBAN KRSIC: [inaudible]

STEVE CONTE: Oh, please.

BOBAN KRSIC: I just saw that we have Norm's name here on two topics. One is here on the left side in security and incident responses and security incident management, and the other one is here on this to be defined topic. So –

UNIDENTIFIED MALE: [inaudible]

BOBAN KRSIC: Perfect. Thank you. Now it's time for a break.

[BREAK]

STEVE CONTE: Okay, so just putting this on the record. We're still on a hard time schedule for Xavier and James, so we actually have 30 minutes. We could either use it to discuss any kind of DAAR related activities or anything else, or just dig up stuff, whatever you guys as a team want to do.

UNIDENTIFIED MALE: [inaudible]

STEVE CONTE: Just for the record, Norm is suggesting we talk about DAAR or domain abuse now until the 4:00 or 4:30 slot. And just reminding people to use mics here so we can get these into the recording.

All right, so that was a yes. Let me pull up and put this into the chat room. There is a DAAR FAQ that was published not too long ago, and that might help navigate the conversation.

I've just put into the chat, into the Adobe room. I notice that not everyone is in the Adobe room, so let me… Let me just put that up. I know. No, I don't.

So anyone who's not in chat, if you want to go to www.icann.org/octo-ssr-daar.

DENISE MICHEL: [Drop it into the Adobe.]

STEVE CONTE: I did drop it in the Adobe chat, but for those who are not – again, ICANN.org/octo-ssr/daar. And that gets you to the webpage, and off of that is a domain abuse activity reporting FAQ. And I'll pause there and let you guys digest that a little bit. Thais is what I'm going to be mostly speaking from, is the FAQ, and I can try to clarify any points on this.

One of the things, Norm, that you mentioned earlier today that… want to address or bring up is you mentioned mitigation on domain abuse. And one of the things that this tool is… So first and foremost, this tool is a tool that ICANN has contracted from an outside party to build. We're using public feeds, different definitions of public feeds. They're all publicly available. There's different cost points to gain some of those feeds.

So again, the various definitions of public depending on how much finances an organization would want to have to recreate this data. The point of this tool was to be able to use this publicly accessible data to show where there could be trending in domain abuse, but not necessarily to point to that trending.

So we want to provide this data in a neutral fashion and then let the community decide whether or not there's abuse taking place. We do plan on having some sort of public reporting off of the DAAR tool. That's going to be subject to any kind of contractual agreements we have with the party that we had built this for us, because we're getting the feeds through the contracted party that built this tool for us. So we might have some inability to post details on specific items based on that contract.

We also have the new GDPR that is in the EU, and we want to make sure that we're compliant with that as well when we post this publicly. Denise?

DENISE MICHEL: I have a follow-up question. Is your team going to provide public information on what it would – when it sorts out what it can and can't publish, what steps could be taken to fully publish – assuming there's going to be some restriction – and how much it would cost? So the ICANN Board and community could make an informed decision on the cost-benefit of securing that information in a way that can be used publicly. That's my first question.

And then I would just note GDPR isn't implemented, of course, until next May. So this isn't currently relevant to that, but of course will be next year after that date. And so are you saying that you're crafting this with the GDPR compliance in mind? I think those are my two questions.

STEVE CONTE: Thank you. Crafting is a big word here. We actually haven't received – to my knowledge and through my conversation with John Crain last week – any request from the community on what type of reporting the community wants to get. I'd have to double check the agenda to confirm, and I will before the day is over, but I believe there is a session in Abu Dhabi about the domain abuse reporting, and that could start the conversation from a community perspective on what they would like to see out of it. And I think part of that discussion will be what we can and cannot display, either contractually or because that data is not

being collected. But then we'll get a sense from what the community – what would be important for them to be able to see on a regular basis. Denise?

DENISE MICHEL: So the community in many ways – and it's well documented – has asked for very specific abuse reporting for several years now. So that certainly is not an impediment. And so I would like to submit an official request. My understanding is that – and it's understandable that ICANN is working with the provider of this data, working through what under the contract can and can't be publicly displayed.

And so my official question is, once they work through that and understand what can't be displayed, will they come back to the public and the Board with the cost of and the mechanism for fully publicizing all of the information, what that would cost and what we could see? That type of thing. I don't think you can answer this now, but I think that is a question that needs to be answered in the context of this work.

STEVE CONTE: Just one single note on that. You said "fully publicizing." But am I correct in assuming that you meant fully publicizing within the limits of the contractual can and can't display based on what we have with contracts with the provider?

DENISE MICHEL: [Every time I say] the contractor for… Yes. No, I'm saying –

STEVE CONTE: Within the limits of what we're contractually allowed to display, even internally or whatever, you're fully displaying is within that subset, not fully of all the data. We won't be able to [hit] all the data, if that's the case.

DENISE MICHEL: No, I understand that. We've been told that this particular contract that enabled, supported ICANN's gathering of this data does not allow it to display all of this data. So I think a logical step from that is, what would a new contract with a new provider potentially, or the same provider, need to look like in order to display more of that particular data publicly, and what would it cost? Is that helpful?

STEVE CONTE: That's something I can go back and ask that question for you. It is helpful, and I don't have an answer, but I'll take that back to John and to Dave. James?

JAMES GANNON: Thanks. So just for some context, I had a conversation – I think it was after Madrid – with John about this. There are about I think ten different feeds that are coming into this, and I know just from data stuff that some of these feeds would never sign up to a full public disclosure because it would destroy their business model.

So I think there would probably be the likelihood that there could be two versions of the tool, one internally with full use, and then potentially a second for community display for feeds that have the ability to be contractually organized to be public, but that it wouldn't be the full breadth of the entire tool.

DENISE MICHEL: Thank you, James. That's useful additional context. That's my understanding as well, and what I'm saying is I think we just need to understand what the options are, is there additional cost involved? I think we need to have more clear and wholesome information on this. So if members of the community or the Board want to invest more time and significantly more money in this, they'll have the information at hand to make that decision.

And just for clarity's sake, I think it's important for this Review Team in this area to understand how this is being structured and what the limitations are.

STEVE CONTE: I agree with that, and one of the takeaways I will do for sure is confirm if there is or is not a session in Abu Dhabi and make sure that that gets back to the plenary. Because if there is a session, Denise, I completely agree that it's something that the review team would probably find interesting at least to attend, if not participate in that dialog there.

Then as far as mitigation, one of the interesting points – I think we can all look around the room and nod that there's been abuse in the DNS

since the Internet became some kind of public entity, once it left schools and the governments and stuff.

And I think one of the things – I know that CCT Review Team has talked about this as well, and I think something to consider too is that now we have a new tool that potentially highlights abuse or at least trending on abuse, but I think in some ways too that the tool is so new – even the abuse has been around for so long, we have such a low data collection point that as far as the mitigation goes, it becomes a more difficult conversation to have right now as it might two or three years from now on how to use the data that's now being collected to mitigate.

And I open that up to the table, because we could use some ideas on that too. But I think without a lot of data behind it – we're showing peaks and valleys right now, but we're not showing really – we don't have enough data to have trending going on that's enough data to be reliable to say that it's a trend. We can talk about what our historical knowledge is on the individual level because we've all seen trends, but without the data to back it up, then it becomes nonneutral opinions and not a neutral analysis based on data.

NORM RITCHIE: Yes. I'll nod as well. Abuse is a pretty broad term as well in and of itself. I think that is not helping the discussion, just [use] the broad term talking about abuse, because you have spam, you have phishing, you have command and controls and botnets, droppers. There are all kinds of different types of abuse that could be occurring, and when you lump it

all together, you kind of lose something there because typically spam overwhelms everything else.

STEVE CONTE: That's a good point, and something I should bring up is that these feeds that we're getting to populate the data within the DAAR tool, they call it reputational data. So the feeds that we're getting are from various organizations such as I think Spamhaus is one of them, or there are other ones out there too. But they're providing what they call a reputation score on that, so I think the abuse that you're mentioning, Norm, is right on spot. I think we're seeing this tool will speak to some abuse that is more easily identified spam, malware, things like that by reputation scores than we might be seeing from a technical abuse of the DNS protocol. And that's just something that – maybe it's not one tool to rule them all. Maybe there'll eventually need to be more than one tool to look at different aspects of DNS abuse and the definition of abuse. Eric?

ERIC OSTERWEIL: Just to sort of pile on a little bit, but to kind of clarify a little bit I think to Norm's point, the types of abuse and the motivators, the incidents change a lot depending on what the actual abuse is. And so conflating them can kind of wind up with metrics that maybe aren't as insightful as they could be.

So what you need to do – there'll be a lot lower incidences of some of the more malicious stuff sometimes, and so breaking them out separately whether – so a reputation score for spam might be

completely different than a reputation score for C2 or stuff like that [infrastructure]. So I think we saw that –

STEVE CONTE: Yes, we can definitely filter out, because you're absolutely right. Spam was a huge indicator on that, and we were able to uncheck spam and bring in different metrics based on that. Because you're absolutely right, spam would just say, "Oh, the whole Internet is bad. We've got to stay away from it."

ERIC OSTERWEIL: Yes, and that was one of the things when people start correlating pricing models with incidence. There are some types of abuse that correlate better with pricing than others, and then if you take sort of the massive dataset and therefore we can conclude something about all types of abuse. You wind up coming up with an improper sort of cause and effect. So keeping them separate until you understand whether there is some sort of underlying commonality – which there may not be – anyway.

STEVE CONTE: Noorul.

NOORUL AMEEN: I understand the DAAR tool is taking some different kinds of feeds. One we were discussing is the [inaudible] framework for the spam kind of issues. So my question is whether you're considering – I understand

ICANN is not in the operation part of the registries, abuse and all the things, but still, you consider domain registration abuse that's like problem domain registrations, those kinds of issues also in the DAAR report too?

STEVE CONTE: That's an interesting question, and one that's harder to quantify through a feed, because if I decided to register a random string of characters on a TLD, is that intent to be malicious?

If we look statistically, especially with some of the DGAs, the domain generation algorithms out there, there's a good chance that, yes, it probably is going to be malicious, but there is no sense that – a random string of characters that might look random to you might be an acronym to me and have actual meaning.

So getting something like that is going to be harder – [I see a queue,] Noorul and James – to quantify at a feed level, and I think would be interesting to hear from the community how we could get trending on that. Noorul and then James.

NOORUL AMEEN: This is a specific parameter if you analyze the configurable infection, the DGA has been generated, the DGA has been used to register problem domains in a larger perspective. In a short span of time if you're observing domains through the random characters, that's a proper feed for a domain abuse.

JAMES GANNON: When you look at domains through the lens of what is abusive and what is not, your statistics can be very skewed. There are massive use cases for DGAs that are not for malicious purposes. And by quantifying a feed or reputation score based on whether it's a DGA or not, that's a very rocky road to go down. I can start listing software and use cases for DGAs which are perfectly legitimate and which are used for general, legitimate business.

So be very careful when we talk about DGAs and using generated domain names as an indicator for abuse or fraudulent activity, because that is a component and it is used within that sphere, but remember that there is a huge set of use cases for DGAs which are legitimate business. And if you're just looking at the string, then there is no way for you to determine one way or the other what that is in a proactive manner.

STEVE CONTE: Before that, just as a reminder, let's keep this conversation – because this is a good conversation but I think it's more of a plenary conversation at that level, let's keep this to the area surrounding the DAAR tool. Norm, did you have a comment?

NORM RITCHIE: No, it would have been [plenary.]

DENISE MICHEL: Can you update us on when we're going to see this and what the current status is?

STEVE CONTE: The current status as far as I know is that, again, I think we're having a session in Abu Dhabi to start determining what the community wants to see out of it. As far as an interactive tool, I don't know what the status is. This is a John or Dave question, and I didn't ask John last week if there's going to be an option for an interactive tool or if we're going to be asking for reports that are common. Because as you mentioned earlier, there are going to be costs somewhere on that, and I think that if there is a session in Abu Dhabi, that might get that dialog going.

As far as the tool itself going on right now, it is active. There are people – because it's contracted, we have a limited seating on that. John Crain and Dave Piscitello obviously have users. We also have given seats over to some members of Compliance and some other teams as well.

So we are actively using it within ICANN for some purposes, but not on a broader level. And because of the seating challenges with contracting, that's why I'm not sure if an interactive report building tool would be fiscally viable at this moment. I think this is something we'd have to look into.

Any other questions about the use of the DAAR tool or anything else around that? Ameen.

NOORUL AMEEN: Once you've got the outputs of the DAAR tool, [let's give a] fine grained level like [inaudible] domains or what kind of domain abuse it has got reported. Abusive domains and abuse domains. On what [inaudible]

called abused, what is actual reason behind that? Let's give further detail.

Say [inaudible] domain, is it [inaudible] used for phishing activities and [inaudible] level details available to DAAR.

STEVE CONTE: I'm not familiar with the registry contracts or Registry Agreements, or the Registrar Agreements. My understanding is that we're starting to look on using the tool internally to look at the trending to see if there are abusive domains that specifically point to a registrar or registry that through our contractual relationship with them, we could take some kind of compliance action.

It's a new tool, we're learning how to use it internally as well. Having a lot of data isn't always a useful thing, so learning how to sift through that data to get good results that are within the bounds of ICANN's mission and within the scope of any kind of contracting is still being a lesson learned right now, and I think as the tool is exercised more internally and we have a better understanding of what those trends actually show us, I think it'll help us as ICANN serve our mission better through enforcement of contracts.

John Crain is on a plane, he's heading to a meeting out in the East Coast, and Dave I think is training I believe in Europe somewhere. So any questions for them I'll have to capture and get back to the subteam on.

NORM RITCHIE: Just one broad question. Is there a plan to make the information of DAAR actionable by the registrars and registry? Currently no, it's just kind of a report, right? They can't really act upon a report.

STEVE CONTE: It's an interesting question. I'm trying to figure out the scope of the question and how to frame an answer. Whether there is a plan or not, I can't answer. Whether or not it'll become a part of [inaudible] policy to use the results of that tool in order to justify or strengthen an action to mitigate abuse – I don't want to say action against a registry or registrar, because it doesn't necessarily need to turn into a legal issue, but using that data to help a registry or registrar mitigate the domain abuse. Because the bottom line of what we want to do is we want to make it a better place out there for the consumers, not necessarily say, "Well, this registrar or registry is a bad actor," and because all these other people are bad, that doesn't necessarily make them the bad actor.

So working with the registry or registrar, using this data to highlight that there's abuse going on within that space, within that label, I think will become eventually some policy within whatever department works with those registry and registrars. Whether or not there's going to be a plan, I'd have to get back to the various departments and see.

NORM RITCHIE: Yes. That was exactly the question. It's not so much to identify bed registrars or registries, but to help them. Yes, exactly that.

NOORUL AMEEN: I understand that the DAAR tool will accept different feeds like DNS logs and abused domain list applied by different vendors. My question as a user perspective, if SSR, we observed a malicious domain is involved in that activity, shall we report to the DAAR providers or something like that?

STEVE CONTE: No, but that's not the answer you're asking. I think from a consumer's perspective – I'm going to use my mom. If my mom noticed something bad going on within a domain, the DAAR Team is probably not the right team to go to. And unfortunately this might be a messaging thing, because consumers on the whole – if my mom ran across a bad domain, she wouldn't know where to go.

And so maybe that's a community messaging more than an ICANN messaging, but I suspect that the consumer would be going to the registry – and this is my opinion – to complain about the domain. The registry might pass that on to the registrar because it falls under them, and the process goes from there.

I don't think that she would report it – and I'll get to you, I notice you, James – to one of the feeds, because the feeds are going off of reputational data, and I'm not sure how they collect and how they define reputation. So I think from a consumer perspective, it's still muddied. The DAAR tool is not going to fix that, and that might be a bigger issue in itself. James, and then a follow-up there.

JAMES GANNON: I think if a CERT has a feed available, they could approach the team, but only if it's in the same format that the tool is expecting, etc. If it's individual domains, I wouldn't think so, but if it's a feed of, "Our CERT looked at this systematically and we assigned a reputation score based on this…"

STEVE CONTE: Absolutely. DAAR would absolutely be open and willing to accept new feeds. But from a consumer perspective, my mom won't be creating her own feed.

NOORUL AMEEN: That's a question, actually from an organizational perspective also. If you see wide abuse rather than a user perspective in our organization, it can be act as a consumer for your DAAR tool. So if you observe anything in common in a wider perspective, it can act as a feed for DAAR tool. That was my point.

STEVE CONTE: Yes, that's a good point. As the SSR team and Office of the CTO, we are open to looking at new feeds. One item I'll mention before I pass it back to James here is that we as the Office of the CTO had an intern over the summer and one of the objects that we gave him to do was to not utilizing ICANN's name, but try to recreate the tool.

It doesn't have to be pretty, doesn't have to have a nice interface, but use the tool because we are saying that these are accessible data. Recreate the tool as much as he could with no budget, basically, to see

and make sure that what we're saying that this is data, that whatever your budget allows, you could reproduce.

And we have the report somewhere on this. It was both a critique of the tool, critique of the process, and also showing his ability to get some of the feeds and difficulties getting some of the other ones. James?

JAMES GANNON: Yes. That actually feeds into my question. So DAAR as it's built now is mostly build on what I would call "commercial" feeds. They're not all for cost, but they are the big, organized feeds. Has ICANN thought about approaching some of the CERT organizations first or any of the other big conglomerations of CERTs to look at a more – let's call it an open source version of the same type of feed reporting to bring something together into a centralized place?

Because there are a lot of feeds out there by various CERTs and other organizations that are not for profit, they're free, but they're all scattered as opposed to centralized.

STEVE CONTE: That's a great question, and I think you actually hit upon the answer somewhat earlier when you were responding to Denise, is that feeds are a revenue source to people. And so we would love to see an open, public, anyone can access feed, and maybe there is.

The CERTs are out there, but maybe depending on what information that feed provides, if there ever is one, we would certainly subscribe to it or make sure that we subscribe to it. But I think for many of the

organizations out there, even nonprofits, that's a business model for them. That's a way and a case to generate revenue.

So I don't want to give you a flat out no because I can't speak for the entire world. If we see one that's out there that's relevant to ICANN's mission, we would certainly subscribe to it.

UNIDENTIFIED MALE: A bit off topic, I think. I really like that idea a lot. I'd like to pursue that with people and talk with that further. I'm just talking now for the Secure Domain Foundation, not myself. Sorry.

STEVE CONTE: Yes. If there's something that you guys know that you can point us to. You all know John and Dave Piscitello, so please point us to the right places if you know something this week here. Send me an e-mail or hit me during a break or whatever.

And if there's something that we can add to it that would add more value to the tool – because one of the things that the tool does is it takes these feeds and normalizes it. Because some people will say one feed is more reputable than the other or whatever, but they're all giving us some different sense of information and some different sense of the Internet, whatever angle and focus they're looking at.

So the more feeds we can get to it, the more we can normalize the data and get a better sense of trending on what abuse is and how it's affecting the DNS system.

BOBAN KRSIC: Steve, one question. There are a lot of services you can buy, there's a service to monitor your TLD's zone for any abusive content. And because registrars have this requirement due to the agreement, and I know two or three providers of such services, what about this? What do you think about it when you have something? To put it open to the community, what's with the competitors?

No, not the feeds. The feeds are the [ones] because there are companies who provide this as a service and they use also the same feeds for their products which they offer for their customers, and the customers are registrars which have the requirement and said, "okay, we don't have the resources to build our own infrastructure and we buy it from the market."

So there is a market for this. You have to pay today if you want to analyze your zone regarding any abusive content to categorize it. They call it domain management or domain monitoring, I don't know. And what about this one?

STEVE CONTE: it's certainly a viable option. One of the things that we're going to stand strong on is to make sure that the data that comes into this tool is reproducible if you have the financial means to do everything. It's not necessarily a cheap project that we're running. It's a significant cost and there are going to significant recurring annual cost on this too to keep this running.

We want to make sure that whatever feed we get, that any consumer organization could also get that same feed. So we don't want to necessarily get any kind of custom feeds, unless that custom feed can be either reproduced for another entity or given out freely, or some level in-between that.

In some ways, I think what you're describing is kind of what's going on already. We're accepting reputational feeds, whereas it sounds like you might be looking – the provider, registry, registrar or whatever might be looking for more specific information not only around their label, but what pieces on their label are going to flag as something that falls within the contractual compliance of the agreement.

We're open to it. Anytime we add another feed, it adds a little bit more money to our annual recurring, and again, we stand strong that we want to make sure that it's reproducible by anyone if they have the right means. And even to a lesser so, if they don't have the fiscal means. Like I said, my intern was able to reproduce some level of the DAAR tool using R programming language and just a down and dirty CLI. James.

JAMES GANNON: Just a quick point of order, I suppose. We're 10 past the hour. Are we holding James and Xavier back?

NEGAR FARZINNIA: They are ready and waiting outside.

UNIDENTIFIED MALE: [What we're doing is sort of] space filler. We can fill other spaces.

NEGAR FARZINNIA: I will call them in now.

STEVE CONTE: So we're temporarily putting – or permanently putting [them out] with the room. Feels like DNS abuse and DAAR on hold for now. We'll revisit it either during another space of time, or at 4:30 we'll just do a sense of the room on whether or not we need to continue this discussion.

So with us, we have our Chief Financial Officer, Xavier Calvez, and James Caulfield, our fairly newly appointed Vice President of Enterprise Risk Management to have a discussion about the topics that they hold within their purview.

I don't know how you guys want to run it, if you want to start with an overview or, Boban, if you want to open.

XAVIER CALVEZ: I was going to suggest if you would like to bring the picture of what you would like we cover and then we'll make sure we speak to that, and then we can iterate with questions.

BOBAN KRSIC: Okay, perfect. Thank you. First of all, I want to thank you for your time and to be here and to talk about the business continuity related to the SSR of the unique identifiers. That's the scope of what we are talking

and the idea is to start with an overview of the processes and to say, "Okay, that is our risk management process as far as the business continuity strategy," and then to go in a dialogue to focus on these topics, which are relayed here in the agenda.

So I would propose start with a process overview with a big picture that helps us to go in detail and break it down.

XAVIER CALVEZ: Thank you. So we'll do that. Thank you, all of you, for being here. You're spending time out of your lives for this review, and this is very important, so the thanks should go to you.

So James will try to provide that perspective. We'll probably want to put risk management as whole into the picture and then see how business continues to fit into that because it's a part of it, so it may help the understanding for everyone. It also may be clarified the scope of what we ultimately want to help you focus on. So I'll let James do that.

Again, he's [inaudible] a few months ago in charge of risk management at ICANN, and as part of that, in charge of coordinating the various activities that contribute to business continuity that may be carried out in different parts of the organization from a responsibility standpoint. But from a coordination standpoint, James has [inaudible], which is why he can speak about it. Thank you.

STEVE CONTE: I'm sorry. Just to let you know, we don't have anybody remote, but we are recording this so I do ask that we stick with mics whenever speaking. Thanks.

JAMES CAULFIELD: Great. Thanks. Got it. So I'm James Caulfield, new here at ICANN, very glad to be here. Just to be clear, we're going to talk about risk management, generally, business continuity.

UNIDENTIFIED MALE: Risk management [inaudible].

JAMES CAULFIELD: Okay, great. Thanks. So I've been here since the end of April and ICANN does have a number of important risk management functions and capabilities in place. We do have and had important things like you would expect in a risk management program, like risk register where there is a process to go through and identify risks in the organization, what those impacts would be and to rate those risks and make sure that those are brought to the attention of management in the Board. So that's all expectations, what I would expect in a mature or maturing risk management context.

I think that what had been in my role in ICANN is to help that move forward. I think it's in a good place now and it's my job to improve that, so I don't know if there's specific questions around that.

DENISE MICHEL: I think really more a question for the record, if you could provide, one of the things this team is, of course, looking at that is required that we look at is implementation of the first Security Review Team recommendations and their impact. And in line with that, and the questions around risk management, could you provide us detail on how risk management has been staffed at ICANN since the SSR1 recommendations were adopted by the Board? Thanks.

XAVIER CALVEZ: So thank you for that question, and your question is more specific in relation to the coordination of the unique identifiers or more generally for the entire of ICANN?

DENISE MICHEL: Yeah, it's specifically related. So we're required to look at the risk management framework as it relates to the SSR remit and mission, so it's whatever you think is most useful in answering.

XAVIER CALVEZ Okay. So in that respect, what I'm going to answer is not going to be just carved specifically relative to SSR or unique identified coordination because the organization, starting to answer your question, the structure of the organization which we have in place is not designed in this fashion.

And the staffing question is fairly easy to answer. We have James in risk management today. When the function was created in 2013, the risk management function was created in 2013 to coordinate the efforts of

risk management. It had one person, then added a second person, but the dedicated resources is a picture that's a little bit misleading because risk management is not just about one person doing something in an office as part of the organization.

Risk management is performed, as you know, throughout the organization by every individual ICANN staff member who has responsibilities to operate a function that may be at a risk. That risk then becomes the "ownership" of that manager who manages that function.

So we have, which I think is a very common model, a model where the responsible [effectivities] or functions are also owners of the risks associated with that function whether it's the risk created by the function or a risk that exists in the environment and that simply is related to a function.

DENISE MICHEL: Yeah, I didn't mean to take up the team's time on this particular issue. It would be great to get this in writing, when you guys have a chance, to staffing for the risk management framework responsibility, understanding how risk management works in a corporation. We're talking about the position that he fills, basically now, of coordination, facilitation, tracking, documentation, and all of that.

And this is just something we can get in writing whenever you guys have time to give it to us. I didn't mean to delay your presentation. Thank you.

XAVIER CALVEZ: Sorry. We will not just speak about James's position because it's "the tip of the iceberg". It's not giving the full understanding of the resources dedicated to the topic. But we'll provide a written answer to that question. Thank you.

JAMES CAULFIELD: Sure. Sorry, [inaudible] business continuity. Okay, so then specific to business continuity, we do have what with generic would be called business continuity plans. I think initially, or the first thing to say is that there is a little bit of different terminology used differently, depending on the situation. I think business continuity or business continuity plan is often used generally for what do you do for your business resiliency which is a number of different steps. And per se, we're not a business. We're an organization that works in that community, so even that is a little bit tricky in terms of terminology.

So I would say in terms of an organization, we certainly have resiliency plans that help us continue to carry out our work. There are different steps in resiliency beginning from an event or an issue in how do you handle the immediate impact. We have plans around that. How do you make the decisions in a crisis? We have that, and we also have disaster recovery plans that help us know what it takes to get back online, especially as it relates to systems and networks.

ERIC OSTERWEIL: Sorry, just one question that may be out of place, but do you guys subscribe to any third party audit frameworks or anything like that for those sorts of [business] plans or anything?

JAMES CAULFIELD: We don't specifically for disaster recovery or business [inaudible]. We do for some other areas, but we don't. I think that amongst the people that we have in the organization, we have enough professional experience in that. In fact, developing a fit for purpose plan is often much better drawing from those kinds of frameworks, and also, being familiar with those taking the best parts and looking at those for creating our own plan.

My personal opinion is that you don't want to shove anything into a cookie cutter framework, so having set up a number of business resiliency plans before and also looking at the organization and a number of things that have been put in place and the professionals that have been hired to manage these processes, I think that's actually important to keep in mind, that we have people who have experience doing these kinds of plans, extensively even.

So to answer your question directly, the answer is no, but I think that, certainly, I should better approach this.

NOORUL AMEEN: Do ICANN have a crisis management plan or something like that?

JAMES CAULFIELD: Yes.

NOORUL AMEEN: Publicly available?

JAMES CAULFIELD: I don't think so. I think, often, one of the issues with crisis management for business planning in general has to do with situations where an organization is vulnerable or an organization is literally under attack, and so to make those kinds of things [inaudible], I think it would not be prudent in any company or organization.

BOBAN KRSIC: How do you identify the critical services for them to have provided a business continuity plan? Because I don't think you have for everything a business continuity plan, only for critical services. And how does the process behind that? And how do you know, as an organization, are the business continuity plans effective? Have you tested them or [regularly] tested them? What was the outcome? Do you provide lessons learned, and so on?

JAMES CAULFIELD: There has been a proper analysis of that, and so, I think that that was covered and there has been for critical systems testing of the [time].

JAMES GANNON: Thanks. Sorry, I missed [to introduce]. By the way, I'm another James.

You may have spoken to this, but just for my own clarification, are you still operating under the contingency and continuity of operations plan with regards to the shared services that is now with PTI? Or has that been rolled up into a bigger business continuity planning document?

XAVIER CALVEZ: Sorry. I'll give a shot at the answer, but in answering, you're going to check if I'm actually speaking to your question. I'll complete that I'm going to say now, later. Most of the business continuity planning activities that we have in place are functionally driven, whether they apply to the rest of ICANN organization or to PTI.

Having said that, there is also, as we know, very specific business continuity planning activities relative to the IANA functions which are carried out today within the PTI entity.

So those that were specific continue to be carried out specifically for PTI by PTI delegated resources as they had been before, and those that were carried out across a given function, if that function supports were services ICANN organization and PTI, it's that function that will continue, carry out its business continuity activities in the way that the function was doing it.

So it's a bit of a status quo from the perspective of before/after PTI. The activities continue to be carried out either functionally by the functions that were carrying them out before if they continue to be servicing both the ICANN organization and PTI. And for those that were specific to PTI, they continue to be specific to PTI.

Now, you can correct me if I haven't answered correctly your question.

JAMES GANNON: So I'm going to give some context and then a follow-up. So first of all, we as a community are aware that there is a CCOP, the Contingency and Continuity of Operations, which was specific to the IANA department previously. Okay?

So for context, we've never seen the contents of that document because it's confidential, etc. So I'm going to extrapolate some questions and I may be incorrect having not seen the contents of that document and Steve is going to come in here and I expect [inaudible].

JAMES CAULFIELD: So just on disclaimer here: as you extrapolate and provide us questions, please remember that any pauses might not be awkward silences, but methods of which we can find the question and still keep it an open dialog.

JAMES GANNON: So back to now the question is, my understanding of the CCOP was that it covered both traditional business continuity management and it covered parts of the IT systems which were operated by the IANA department. This may or may not be correct. I don't need you to give me details on that because that is part of the document.

So with the movement of the IANA department out to a separate entity, a lot of the IT systems are now provided by the shared services

agreement and I would like to know whether an impact analysis or a risk analysis was taken place with the movement of PTI out to see was there any update needed to that piece of business continuity planning that is specific to now PTI as a result of the IANA transition.

XAVIER CALVEZ: Thank you for clarifying. So first, I will want to take this specific topic and your specific question back to the IANA Functions Team that currently continues to operate, that the business continuity [inaudible] in money during to make sure that what I'm going to say is accurate or adequate.

If you think about it in a general fashion, the construct of PTI has had little operational impact on the daily activities that supports the IANA functions. And you mentioned an SLA, for example, that SLA has not obviously changed the activities that were carried out, but simply documented them so that the construct of a new entity in which those functions are carried out is documented.

The reality of what was called before the IANA Team that's now part of PTI, what they do on a daily basis has very little change, if sometimes not changed. And the business continuity activities carried out have also not changed.

Where before the creation of PTI, the IANA Team was relying on the Department of IT for support, it probably continues to do so. I know there's I believe one activity of software, [inaudible], that's specific to the support of the IANA functions for those software that are only used by IANA that I think that activity has been a bit more carved out than

others that are more shared. But I actually don't think it was just due to the creation of PTI. I think it was just due to the logic operationally of making that happen.

So I think that the operational reality has changed little. I think that the impact analysis of "What does that change to be in PTI versus not before?" would have been driven due the construct of PTI's legal entity with now-existing contracts that document what was happening. Does that change the operations of the IANA function? Because that's really what the change is. It's a structure rather than different activities.

So that, I would need to check. Have we performed a very specific risk analysis of having contracts versus not having contracts? Does that create a change?

The other aspect is there are a few specific additional activities that have been created as a result of the PTI construct, RZERC, and the customer service committee reviews [and some]. And those have been developed in or continue to be developed, if you think about it because we're still in the first year of carrying them out, with an understanding of how they should be operationalized. And I don't believe we have done a very specific risk analysis of those activities other than trying to design them, if you see what I'm saying and ensuring that they're carried out according to the proof proposals of the CWG on transition.

ERIC OSTERWEIL: Okay, this serves as a follow-on question that maybe you answered and I missed it. So in preparation for any kind of a risk analysis, is there any kind of documentation, either internal or external, about the processes

that are now required to coordinate what was previously done in a single organization, now across an organizational boundary?

In other words, even if everything is working exactly the same as it was, now that there's two organizations, there's inherently more processes, exactly to your point, to bisect the operational duty. So it seems like that might be something you'd do in preparation of a risk analysis. I just wonder if you have any thought or if there's been any work on that, if that even made sense?

JAMES CAULFIELD: Just to speak a little historically, the IANA function even before PTI, it used to be under contract with the U.S. government, at least a long time ago, we used to make sure that the IANA function was modular because if ICANN ever lost the contract with NTIA, they would have to move that to a different organization.

In my opinion, in some ways, the fact that we're now contractually bound with the IANA function through PTI is actually strengthening. And please jump in on that because I don't know, of modern times, but through history, it was more ethereal, the support and process whereas now with contracts, there's specific points that both sides have to adhere to.

ERIC OSTERWEIL: Yeah, thank you. That's exactly to the point I was asking about. Those are the things that I think we would just be interested in understanding, the points at which there is sort of a procedural communication where

before it was sort of under the same auspice, but was sort of segregated logically. That's the point at which a resiliency plan needs to be sure that if there is a communication required between two organizations and one organization undergoes an event, that it's documented what the sort of resilience plan for that is. That was kind of my point. We would be curious about that sort of stuff.

JAMES GANNON: I had three questions. Now I have two because that was actually one of them. One of the other questions was has the risk management framework been updated to ensure that, for example, PTI employees are empowered, will be the word I'll use, to be risk owners with ICANN's risk management framework to ensure, thus, where PTI services risks that then need to get passed over to ICANN because let's say it's something to do with the shared services agreement. Has the risk management framework been updated to account for those, again, organizational differences or is that even needed? And then I have another follow-up after that.

XAVIER CALVEZ: Let James add to that. But I'm glad you added the second part of your question "Is that really needed?" because I don't think that from a modeling standpoint, the change of structure really creates any barrier or created any barrier for that communication to be happening because the model that we used of risk identification and update of the risk register that James mentioned earlier was all-encompassing.

But to the point that I think you were making earlier, before PTI, Elise Gerich's title and role was IANA Department Manager, without going into detail of the title. Now she's President of PTI. Have our processes been updated to reflect that change of title, maybe the change of authority or documentation of authority? They have, but not necessarily as a result of a specific risk identification exercise, but simply as a part of the implementation of the transition recommendations, if you see what I'm saying.

Having said that, to the point that I think you are making, is have we purposefully analyzed the potential unwanted impact of the change of constructs? I think that's definitely a good thing to ensure that we have done correctly, and maybe to document more than we potentially have had. I think that's a good thing to investigate and develop further. Thank you.

JAMES GANNON: One more. My last one, don't worry. So going back to the CCOP, the PTI-specific disaster recovery planning, first of all, I assume you do tabletop exercises for business continuity within the CCOP or the ones that are now specific to PTI, do they occur at a more regular frequency and to a deeper level than the standard for potentially other departments? This is the testing specifically, or the exercises, whichever way you verify that your business continuity is actionable rather than just a policy or a procedure.

XAVIER CALVEZ: I actually don't know the answer to that question. Do you?

JAMES CAULFIELD: I don't know for the CCOPs, specifically. For ICANN overall, yeah, we have. But whether that's been done on the CCOP level specifically, I don't know.

XAVIER CALVEZ: I think it has been done, but the question about the frequency, I don't know the answer, or the periodicity, so we'll follow-up on that. Thank you.

DENISE MICHEL: Thank you so much for taking the time off from what I know is your super-busy schedule to come address these. I've put some additional questions in writing that eventually will make their way to you, but could you give us a sense of how frequently and to what extent the Board is updated on the risk management framework and activities related to this? And also, give us a little more information about the sort of [inaudible] and community participation that's involved in the risk management framework development.

JAMES GANNON: And could I just actually tag something on to that exact question? How do you do risk reporting? So you do risk election, but how is risk reporting in general handled? Is it a C-level forum that's held once a month? Does it go straight to the Board? What is the process around that?

DENISE MICHEL: And if I can just elaborate a little bit more on my initial question, which [inaudible] two questions, one is about the Board and then one is about community engagement and events, and specifically relating to the Board, there is, of course, a Board Risk Committee and so I'm particularly interested in the briefings to this committee on risk assessment and, of course the proposed mitigation measures since 2013 in particular, and what work has arisen following those briefings.

XAVIER CALVEZ: As Denise mentioned, there is a Risk Committee of the Board that is delegated by the Board, the duties of oversight of risk management of the organization so that all the Board members, not just the Risk Committee members, can exercise their duty of care, which includes, among many other things, to ensure that the organization has the ability to identify, evaluate, and mitigate risks.

So the risk register – I'll use the specific process that we have there in place. The risk register is resulting from a process of identification of risks that was put in place at its inception, and then continues to be updated on an ongoing basis.

The risks identified have been, at the inception, the result of the Board input, organization input, and community input as well. To those of you present or participating in the community four years ago now, three or four years ago, there was an exercise where the SOs and AC organizations and their leaders were asked to provide their views of potential risks impacting ICANN. And that input collected from a lot of

organizations made its way into that risk register along with compiling the input from the Board, and of course, there was a lot over that as well as the input from the organization.

So that's the risk register and the risk register is updated on a quarterly basis, so calendar quarter periodicity for evaluating or adjusting the likelihood of the risk occurring in the potential impact of the risk if it would occur. And of course, what are the controls in place or activities that mitigate those risks or that are aimed at mitigating those risks?

That is an update on a quarterly basis that the Risk Committee of the Board receives from the risk management function. It receives what are the main changes to the risk register and what were those changes driven by.

I think we said it before in Johannesburg in a different context, but the risk register is each of the risks of the risk register or owned from an organizational standpoint by someone in the organization who is close to the function that either manages the risk or is in the scope of the risk.

And that information is consolidated on a quarterly basis and provided to the Risk Committee of the Board. That's the ongoing process for the risk register.

In addition to that, the Board Risk Committee has organized on the at least annual basis, sometimes semi-annual basis, a session with the entire Board where the Risk Committee of the Board presents and discusses with the entire Board the risks that are being considered in that risk register with a focus on, depends on the period, the top 10, top 5, top 13, depends on how that came about.

And so that, again, the entire Board has visibility on where the risks are by the organization and what is being performed by both the Risk Committee and the organization to identify, update the assessment and continue to mitigate those risks. So that has been happening either every six months or every year.

The Risk Committee of the Board meets whether it is on the phone or face-to-face between, I would say, three and six times per year. The session on the risk register with the entire Board has happened every year so far, and is actually scheduled to happen in Abu Dhabi. The advantage of the General Assembly meeting is that we have an opportunity to do that exercise with both the departing Board members who have the background and experience of having looked at that during their tenure, and the incoming Board members who then benefit from the history that their colleagues are sharing. So this is a very useful discussion.

And lastly, the Risk Committee of the Board organizes what is called a Risk Workshop. It's not the Board Risk Committee meeting. It's open to the entire Board, for those who want to participate. It's also open to any of the Executive Team members, and it's an opportunity outside of the formality of a Risk Committee to allow for interactive discussions for education and information, both ways, I would say, between more knowledgeable Risk Committee members versus less knowledgeable, or some have a special "team" or expertise. It's also the opportunity for the Risk Committee members and the entire Board members to ask questions to the executives who are present. It's also the opportunity for the executives to provide input, information, or sometimes

questions, also to the Board members who have expertise on their role of oversight. So it's an interactive session.

So that's a summary to try to answer your question, what's the involvement of the Board Risk Committee? It has these three types of interactions. I'm not talking about the detailed working plan that the Board Risk Committee has, but I'm mentioning the meetings that around the risk register.

Then, the second part of your question, which is how much of that then is being discussed and presented publicly. When I say "publicly," I mean both to the community and outside of the community.

So over the past year and a half, we have been discussing with the Board Risk Committee that very point. How do we share with the community and the public what ICANN does relative to risk management?

We are in the process of designing that aspect because that's not something that ICANN has done a lot before, and fairly logically, without having an established risk management function within the organization in the past, that specific activity has not necessarily been designed, and structured, and carried out.

So we have discussed with the Risk Committee of the Board and I have suggested a model of interaction relative to risks that uses the example that some of you know of the Budget Working Group that has also been carried out over the past few years, that has offered the benefits of an interactive discussion, broad participation in the sense of completely open participation, and allowing for interactive discussion.

So I put that on a hiatus over the past couple meetings because we didn't manage to design this process and what content could be discussed during this process until James was here, so this is something that we need to be able to resume.

The big question, which I know everyone has on their mind is, how much can we share and how much can we not share? And, of course, the question is how much do you increase the risks or not by sharing information about those risks or how you manage those risks?

Having said that, there is also, certainly, out of what the organization does in terms of risk management, there should be things that we can share and we can discuss. So I think that that model needs yet to be defined.

We have in mind the model of interaction. We need to design the content that goes through that model of interaction and I definitely expect that it will be quickly interactive in the sense that the first time we have this meeting, I think we'll all lay out those questions of what can be shared, what cannot be shared, what can be shared under what type of conditions, and so on.

So you have been exposed to the question of a nondisclosure agreement, how can we use that type of mechanism possibly to share more information with [inaudible] than the public? So that's in the design, but that's definitely something that the Board, the Risk Committee has tasked the organization to develop, and design, and kickstart. It's not yet been done, to be very clear.

Yes?

DENISE MICHEL: Question. Is there a final DNS Risk Assessment document? Because we've been given something that's marked "draft" and dated May 28, 2014. And feel free to e-mail us if there is because it's a bit confusing about where the final document is or if you do have a final document or whether there is any update since 2014 because all we have is that May 28, 2014 document. And thank you for your full explanation. That's really helpful.

I apologize, and do let me know if I'm being repetitive because I had to step out for a minute, but it seems that the SSR, the first Security Review Team final report described the need for a much more structured process for identifying near and long-term risk, particularly as they relate to the Internet's unique identifiers, and that ICANN should not only publish more information about these risks, but with the understanding that some of it may be sensitive.

It seems that the SSR and OCTO teams do publish, occasionally, important materials that go to the near and long-term risk in this arena. I'd like to have a better understanding of the mechanism for feeding this information into the risk management framework, and into ICANN's strategic plan, and budget priorities. And I think a sub-part of that question [than Dave] is to understand how your position specifically relates to risks in the area of the Internet's unique identifiers, that intersection there.

XAVIER CALVEZ: So we'll look for the final document.

DENISE MICHEL: Feel free, and I should just ask Boban how we're doing on time. And feel free to answer any of this in writing. I don't want to take up too much of your time here.

NEGAR FARZINNIA: If I may, James and Xavier have a hard stop at noon. We have about eight minutes left. I think any follow-up questions, obviously until 12 are absolutely fine, and beyond that, we are happy to take them back and forward a response later.

DENISE MICHEL: And please feel free to just answer this in an e-mail to the team. Other people may have questions as we only have eight minutes left.

XAVIER CALVEZ: Thank you. Coincidentally, the meeting that we have at 12 is with the Board Risk Committee, their Chairs.

So the function of risk management at ICANN is enterprise-wide, organization-wide, and that should actually give you an idea about the specific SSR Review scope and is not specific or distinguished from the risk management of the unique identifiers, which may be either a good or a bad thing. I think that's probably a useful area for you to investigate and think about.

So that function today is “holistic”. You use the example of the information that the CTU office and Security Team office teams produce on an ongoing basis. That is an input that along with the input from other functions, comes into that risk register process of identification and update and will receive the same evaluation as to whether it’s an enterprise-wide risk or very specific functional risk, and whether this is a risk that we want to track at the risk management level of the organization.

I’m introducing here the idea that there may be risks that are important and monitored within a function that may or may not be warranting tracking at the organizational level because you don’t want to and you shouldn’t track [inaudible] risks at the organizational level. You should focus on what those risks are the most important or impactful.

For each function, including the CTU’s office, the aim of those that are specific to DNS security or operational processes, and I know you have other meetings that will come up with, but you’ll be discussing more specific operational processes relative to the domain name system management and contracted parties.

Those from the CTU office or other areas of the organization are received on an ongoing basis as part of this process of input and consolidation of the update from the risk register. But not necessarily as a specific, or tailored, or customized process for the DNS unique identifiers.

I think James has a question.

JAMES GANNON: Just very quickly, just so I understand right, so within your risk classification or scoring, whichever way you do it, there is no, I suppose, multiplier for the potential impact on the unique identifier system. You just roll it up into your standard risk classification.

XAVIER CALVEZ: Correct.

BOBAN KRSIC: [Inaudible] as we have only five minutes left and we didn't touch business continuity and stuff like that, we were talking about risk management.

And please correct me if I'm wrong. How I understand it is that we have certain services and equipment and, as you know, power supply which is shared within different functions. You mentioned that risk is recognized by functional requirement, which is correct. But how do business continuity plans and strategy fit into different functions? For example, who declares disaster in the case that something happens? That's just an example. I am looking for a little bit further view.

JAMES CAULFIELD: Well, our plans do cover that, and so, there is both the capability at the ICANN level and at the CTI level to declare an emergency and manage that through the crisis management process.

XAVIER CALVEZ: I think your point is also about the functional versus transversal part. So, to answer that question correctly, the risks that we monitor are not categorized by function. They are "what they are." They are then owned by someone in a department or a function, so I don't think we have a siloed approach from that perspective.

For anyone to then identify a disaster or a disruption of service, it's going to be whoever is the manager, whomever manages that service. So a lot of those may fit within our IT department, for example. And the functions or systems or services that the IT department manages go across the entire organization, support the organization. So it may come from the user feedback or from the actual monitoring of the services that are being provided.

But I'm not sure I fully answered your question, so can you react to my answer to see how much of it I've answered or not answered?

UNIDENTIFIED MALE: We do have the VRC call and I think we're coming back tomorrow, so maybe we start off with that.

STEVE CONTE: Let's do a little housekeeping here. So I have you guys scheduled again tomorrow. Yes, tomorrow at 3:30. So I guess I asked the team to think about questions. If we could send them, if you have them or if you have the opportunity to send them to the subteam list, that would be great and we can get that over to Xavier and James before they come back tomorrow and expedite, make their time here tomorrow.

Very well, Denise?

DENISE MICHEL: Yeah, I sent some related questions to the full team list.

BOBAN KRSIC: Thank you both for coming today and we'll see you again tomorrow.

STEVE CONTE: So with that, we're at lunch from 12:00 to 1:00, and then after lunch, we have some members of IT which actually might feed into some of Zarko's questions as well.

Is there any reflections, thoughts, items that we need to capture before we break for lunch at this point, that's still haunting people's heads?

NEGAR FARZINNIA: [inaudible] after lunch today, we actually have GDD staff coming in. You're looking at tomorrow's schedule.

STEVE CONTE: I got to switch back to the 9th. Yeah, so we have GDD after lunch. Thanks.

UNIDENTIFIED MALE: Yeah, business continuity isn't my forte, but listening to that, one thing that is not clear to me is the IANA group, which is now called PTI, is that

a distinct entity now with distinct infrastructure, and networks, or is it meshed in with the rest of ICANN?

JAMES GANNON: So PTI is now a distinct legal entity, same staff. The infrastructure, networks, everything are provided by ICANN under a shared services agreement, so PTI is an affiliate of ICANN. So it's a wholly owned subsidiary, basically, and they are provided with the same services that they used to have but now under contract from ICANN. The networks are actually shared as well.

UNIDENTIFIED MALE: Okay, got it.

STEVE CONTE: James, just to confirm. You were part of the transition group. That's why you are more informed than I am on this. Okay.

This might be a question for tomorrow with IT. Same networks as a large term as well because we do have services, root servers are not on this network. So there is segmentation, be it switched or physical, so there might be parts of PTI or the IANA function that don't touch the traditional ICANN flat network. But when they log in, in the morning, they're on the ICANN secure network. From an end user perspective, they're there. And then the servers and the services might be on a segregated network from there. But I'll have to confirm that tomorrow.

So with that, unless there is any further comments or questions, we break for lunch. You're looking at me dubiously, Eric. Is there anything?

ERIC OSTERWEIL: No [inaudible].

STEVE CONTE: Okay. Eric's going to take another Uber excursion, I think.

All right, so we're going to go ahead and pause the recording or stop or whatever works for everyone, and we'll reconvene at 1:00. Otherwise, they're screwed right behind half the table and right in front of the other half. Thanks.

[BREAK]

UNIDENTIFIED MALE: Okay. This is the after lunch of day 1 and if you'll pardon the pause or bring members of GDD in and we can start the next session, which is on Introduction to ICANN's GDD Operations, Processes, Services in Relation to Security, Stability of the DNS. And with us, we have Eleeza Agopian. We have Christine Willett. And we have Francisco Arias looking for a place to sit.

So it looks like with the exception of Brian who's going to come in a little bit later, we have you guys for a nice chunk of time. Boban, I'm not sure

how you want to start this off. Will I pass it over to you and see which one to do?

BOBAN KRSIC: Thank you. So welcome back. Session #2. And yeah, thank you for the time to be with you for the next two hours or at most we have two hours, one hour – I don't know, two hours – and to talk about GDD operation-related tasks. And we have prepared topics on this and it would be great if you could give us an overview about that and to help us to understand this better and to take it in the context of ICANN, our mandate.

We will go into a discussion when we are to interact – it wouldn't be a monologue to us. So we try to figure out how we can or what's the recommendations we can give you how to make this, all things better. So I would say let's start.

CHRISTINE WILLETT: My name is Christine Willett. I'm Vice President of Operations for the Global Domains Division here at ICANN. And in that role, I'm responsible for service delivery to largely to our contracted parties. And I'm responsible for a number of functions and capabilities being operational on behalf of contracted parties and rights holders so that contracted parties can fulfill their obligations. Things like our Trademark Clearinghouse, the Centralized Zone Data System. We administer and coordinate some of those functions. And when I say service delivery to contracted parties, you may be familiar with the RSEP process. But by and large, the operational processing are a series of contract

management and contract administration function that we do on behalf of our 3000+ contracted parties. We provide these contract parties of course our 1200 Registry operators, nearly 3000 registrars today.

So that's the basis of operations. But I see based on the agenda there, the topics you also have included number of what we would call technical services. So I'll turn you over to my colleague, Francisco.

FRANCISCO ARIAS: Can you hear me? Okay. This is Francisco Arias, Senior Director of Technical Services in GDD. I saw some questions in the Excel workbook that contained agenda. I was wondering how do you want to proceed on these topics? Perhaps we just open the [first] four questions or – some of the topics if I recall like the first ones that are here and CZDS and SLA Monitoring System do not have the specific questions. So I wonder if you have any specific questions that you want us to cover.

I have to note that on the SLA Monitoring System, we have had previous conversations in Johannesburg and also you sent me some questions by e-mail and I provided response just a few days ago. Do you have any other questions on these topics? CZDS o SLA Monitoring System?

JAMES GANNON: Not on the SLA Monitoring System but we might start the discussion about EBERO if that's okay with everybody else. Yeah.

STEVE CONTE: Perhaps Francisco can do a briefer overview on what the different products are that you're talking about. And then maybe these questions or discussion can spawn from that just to make sure that everyone's on the same page at the table here. I know that you spoke to us about SLAM, SLA Monitoring in Johannesburg. I know Norm's not in [this] community but Norm's in the Review Team and was in Johannesburg so maybe even a small snippet of what's going on there too and then have the discussion grow from that point, if that works for everybody.

FRANCISCO ARIAS: So the first system listed in the agenda is the Centralized Zone Data Service, CZDS. That's service that was envisioned as part of the New gTLD Program. And there was a group formed with members of the community to design a service that could improve the situation thinking that in the [layers of] TLDs, some of us were available for each of the different [layers of] TLDs. However the requests have to be sent to each registry and obtain, download it from them for each of the different registry systems independently.

So the group that worked on that from the community lecture, they came up with recommendations that had to have a centralized place where users interested in having access to the zone files from each of their new TLDs can go to one place where they will have a user and from there, they could request access to any zone file from any of the new TLDs. And as for the registry side, the registry can go to one place where they could approve or reject according to the conditions described in the Registry Agreement that request for access to the zone files.

And so that's the system we call the CZDS. That's the system that ICANN operates and that system serves that purpose. And so that's CZDS. On our system, we manage – sorry, I don't manage CZDS. But the old system that is in the list is the SLA Monitoring Systems, that is something that my team do manage, and that system is concerned with monitoring the performance of the four services particularly in the Registry Agreement that is DNS, DNSSEC, WHOIS or I should say RDS, and EPP.

Those four services are describing the SLA Monitoring Agreement of the Registry Agreement that's specification then. And then there are series of service level requirements that I can find for each of these services. For example, the maximum response time, a maximum loan time in a monthly basis. So this system, the SLAMers would call it, is monitoring this on an ongoing basis in the case of DNS. The recordings are done every minute to each of the names that were submitted, new TLDs. And as a matter of fact, we monitor all the TLDs, only layer ccTLDS also. So ccTLDs and the root zone. Of course, we don't have SLAs with ccTLDs or the root zone but we nevertheless monitor it so that we know what the status of the service in the case of DNS.

We also monitor RDS. RDS being at this moment WHOIS and [inaudible] and hopefully [3 and 480] and hopefully eventually would be also RDAP or WHOIS will be replaced by [RDS] I should say.

We are not currently monitoring EPP. That's officially to be added to the system. And now, the SLAM also helps in coming to your question or the topic you were raising, James, about EBERO. SLAM or the SLA for sublevel agreement for new TLDs has one section, section 6 that defines

what are called emergency thresholds. These are thresholds that are about the SLA, for example, the allowed downtime for WHOIS is I believe around 64 minutes or something like that per month. And above that at 24 hours long time it's the emergency threshold for RDS.

So if our service reaches that emergency threshold then ICANN has per day either [clause] in the Registry Agreement as the ability to take control of that TLD that is failing in one of these four services that I mentioned plus that [inaudible] and we can take that TLD and put it in let's say in the emergency maintenance by ICANN. Of course, this was assigned as a way to mitigate the risk of having a TLD that is failing, leaving the users so that it will be [raised] on the end users. We have a possibility of using the services under those TLDs that will be affected by this.

We have provided data in the DNS symposium that Eric organized in May in Madrid. We [brought] the data about the things we have seen in the SLA Monitoring System. And recently, we shared a little bit more data at the request of the Work Track 4 of the New gTLD Subsequent Procedures so there is also more public information that has been provided regarding the SLA Monitoring System.

We are also planning to – we have a break in our roadmap to have a public webpage where we intend to publish some data about what we see in the SLA Monitoring System. To be clear, we don't intend to publish a specific case. We don't intend to name the [part] that will have failed but give more high-level view, for example, the number of cases where the emergency threshold has been reached or some percentage that we think that are relevant to them, maybe relevant to

the public. But that's something that is still not available. We have a technical limitation in the current system that does not allow us to do that. That's our high-level review of those two systems and our relation to EBERO in the case of the SLA Monitoring System.

DENISE MICHEL: Thank you, Francisco, you covered a lot in just a little bit of time. So you mentioned the May DNS Symposium. Could you dig a little deeper for us on the reported 32 out of 37 failure, RSP failures. Perhaps to the uninitiated, that seems like a very high percentage of failures. Was the ICANN staff surprised by that? Did it cause you to change any of your activities or strategies in relating to this service? I have those questions off the top of my head.

And then you noted that you're contemplating the public webpage. If you have an ETA on when that would be up, that would be useful to have. And if you could also provide a little explanation as to why the failures need to remain secret and why the public should not know which one failed this test? Those are my four questions. Thank you.

FRANCISCO ARIAS: Let me see if I remember. So around the ETA, I'm afraid I don't have an ETA at the moment. And then you mentioned 32 cases that have reached an emergency… Sorry let me turn this off.

DENISE MICHEL: It's reported that 32 out of 37 had failures.

FRANCISCO ARIAS: Yeah. So there's another 32 there that the one we reported to the Work Track 4. There had been 32 cases that have reached [inaudible] case. Okay, yes. So, 52 RSPs out of 57 that we have identified into ICANN since the [registrar] have the obligation to identify the RSP have had at least one DNS service failure, not an emergency threshold, which the number for the RSPs that have reached an emergency threshold is 11. And I just wanted to make that differentiation.

A DNS [inaudible] can be very short and we have five cases in which the services failed for a few minutes. There appears to be some [inaudible] top that we have seen in some specific implementation of DNSSEC that [leaves] the zone files trailing DNSSEC for just a bit of minutes and I think that we understand that whoever maintains it is working to fix that. That's something that has been reported to us when we asked them to raise as to what happened there and that's what they have told us.

So that's an example of when the small [leave] failures that we have seen in DNS. And that's contrasted to cases where you see an extended failure that reaches the emergency threshold. Those are the 32 cases that we have seen so far. And I forgot –

DENISE MICHEL: I'm sorry. I'm getting a little confused. The 32 cases are the short duration, not the emergencies? So the 32 are the emergency?

FRANCISCO ARIAS: So 32 are the cases that reach emergency threshold.

DENISE MICHEL: Yeah. And so can you give us some more context for that? Just to repeat my question. Is that surprising to your staff? Is that in line with what you expected to happen, if it was a surprise or generated concern on your staff? Did it cause you to take a different course of action to prevent that from occurring again? Can you give us a little more context of –

FRANCISCO ARIAS: I don't know if we had expectation what will happen. What I can say is we were wondering what's going to happen and we prepared for the worst. I think our operations work. We didn't have more events than we could handle. That has not happened. So in that regard, I think we are within parameters of what we can handle.

DENISE MICHEL: [inaudible]

FRANCISCO ARIAS: Right. Right. In terms of what we have seen, we have raised this topic with the community and particularly with the Registries Stakeholder Group. There has been at least one session that we have with the Registries Stakeholder Group in Johannesburg if memory serves where we discussed this topic. There is a discussion group within the Registries Stakeholder Group that is looking at this issue to see what can be done, what can be done better so that we don't have this.

What can we do to mitigate the occurrence of these cases so we have less emergency threshold cases. And there is also discussion within the New gTLD Subsequent Procedures PDP where they are talking about on a more long-term basis. The discussion within the Registries Stakeholder Group as I understand is focused more on what can we do given the current contractual and policy requirements and probably more of a short-term view of what can be immediately improved while the discussion in the New gTLD Subsequent Procedures I understand is more long-term view and multi-purpose situation in the long term perhaps in the next round, I guess.

DENISE MICHEL: So for those of us who are the last initiated in this area, 32 out of 37 failures seems to be very high and very concerning. But what I'm hearing is that it's not and it hasn't caused a change in the activities. And that I guess in line with that then the 11 that reached emergency levels, is that also something that is apparently within a normal or acceptable range given the number that you have? And again also does it [inaudible] your system and so you don't think it requires any change in activities, metrics, that type of thing?

FRANCISCO ARIAS: Right. So, a few clarifications there. In terms of the cases that has been – 32 cases have reached an emergency threshold. So that's one number. The other number is 32 out of 57 RSPs. This happens to be the same number – have had at least one DNS service failure, not an emergency

threshold failure. And 11 out of 37 RSPs have had an emergency threshold. So just to define the numbers there.

Regarding whether we have seen the need to do something, I think I discovered that if we're [initially] talking with the Registries Stakeholder Group on what can be done to improve the situation immediately, and that was also picked up by the New gTLD Subsequent Procedures discussion. So in that sense, that's what we have done so far besides – of course, [exercise] in the necessary [inaudible] we are ready in case we need to do an EBERO transition.

JAMES GANNON: So the emergency threshold breaches are something of great interest to me. Inasmuch detail as you feel comfortable saying in a public sphere, I'd be really interested in you walking us through the process. So your department gets notified by SLAM that an emergency threshold has been breached.

Can you walk me through what happens after that so I understand there is meetings called with the Registry Service Provider, and then you walk through whether to initiate an EBERO or not. Can you walk us through in as much detail as you can that process and what documentation or procedures you have around how you manage that going from SLAM notifying you that an emergency threshold had been reached to making the decision whether to initiate an EBERO or not.

FRANCISCO ARIAS: So in terms of documentation, there is two things that immediately come to mind. The SLA itself describes what needs to happen if there is an issue and how things are to be measured, etc. There is also the registry transition process, which is a process that is included by reference in the Registry Agreement and that describes how among other things, the EBERO transition is to happen at a high level.

We of course have [inaudible] procedures internally that I'm not sure we have published. Now in terms of how things work in case of an issue, you're right. Things start with the SLAM System. If it detects that service is down, then we'll get an alert. And its value depends on the service. If it's DNS, we'll get an alert immediately at the moment the service fails.

It's not really immediately because in order to avoid false positives, it's really three continuous testing periods which is three minutes of failure, that's when we receive the alert. And as to the rate that receives an alert, we are saying that the service is down, etc. And in the case of RDS, the first alert comes at 10% of emergency threshold. As a matter of fact, from then it's the same for DNS and RDS. So 10, 25, 50, 75, and 100% would receive alerts in each of those cases.

The alerts come as e-mails and also as phone calls. In the case of the phone calls to us the ICANN staff, they come and we have two sources just in case to avoid one of the two sources failing so we have automatic calls, a robo calls us, it's called… So it's this thing calling us saying, "You have something, go check it." Or we also have semi 24 knock and the function there is outsourced and they also call us, so we have a human looking at something that has been calling us and saying there is something to be done.

And that's when my team – we have an on-call for all, we rotate who is first and then we have a rotation so in case the first person doesn't respond to the call then the next, and so on and so forth. So when we are woke up and we need to check this issue, we start trying to find what's going on if we can.

As much as we can see from outside what is potentially causing this issue. Well, I guess as you say not what is causing the issue but what we are seeing, what are the symptoms. And then we reach out to the registry. The system also reach out to the registry but it's just to say it's the same kind of alert we received saying that there is a failure. It doesn't give you the details.

Right. So that's where we start the communication with them by e-mail and by phone to see what we are seeing. Give them more [details] to try to [inaudible] to [resolve] the issue. And so that process goes on until the issue is fixed. Most of the cases are fixed by the registry before we reach the emergency threshold so that there's nothing else to do. From my team of course then Compliance enters and they follow the process and there is a request for root cause analysis and what they are fixing in order to avoid this from happening at that time, getting into the compliance land here.

Now, in the case where we reach the emergency threshold for the internal procedures, we start preparing before we reach emergency threshold. We have an internal team that is beyond my team that we have the [inaudible] ops also involving that process. A [inaudible] team and so we start to assess the situation and see what is behind this TLD to assess who will need to call an EBERO and what would be the impact.

So things that are coming to the process to decide if there is a need for an EBERO then are things like are they registered? If it's a TLD that has not launched, then what's the point? And the type of failure. Like for example, is this only DNSSEC failure? Which by the way is most of the failures in what we've seen the DNS service. Because just through DNSSEC failing [inaudible] the DNSSEC not responding. And then in those cases, we look at do they have [fine] domain names in the root TLD? If they don't, well then we need to consider what are the implications of taking it?

We also take into account the response from the registry. We are talking with the registry and in all the cases, we have always been able to reach the registry and have conversations about what's going on. If we have a sort of situation as the registry knowing what's going on, knowing what's the problem that they need to fix, and we have a reason to believe that they are going to fix it within a certain timeframe, we need to consider that timeframe that they have to fix that issue. Again the time it takes to do a newer transition. Newer transitions are not immediate.

JAMES GANNON: Sorry. I'm losing the train of thought. So this process of assessment, is that judgment call by the team or is this something that's formalized and proceduralized? You give between 100% of the threshold being breached to 200% to get a response and between 200% and 300% you consider initiating EBERO. Is that a process that is formal? Or is this people sitting around a table and making decisions? How is that handled?

FRANCISCO ARIAS: No, there is a procedure. We have internal procedure to deal with these cases. And [except] decisions or information is being analyzed before we reach the 100% of emergency threshold. But there is also a component of adjustment. We have to consider, for example, the registry have a good grasp of what's the issue. Have they identified the problem? Can they fix the problem? And so on and so forth. So that there is by necessity so far. No need to have some judgment internally besides the process that we have at hand.

ERIC OSTERWEIL: So I have a handful of things that sort of occurred to me as we were going on so I'll try to string them out instead of [inaudible]. So what's the nature of the infrastructure that you use or you plan to use for EBERO? In other words, if you got – and maybe this is sort of off-topic for this discussion, in which case we can punt it to other places – but I think in the event that the organization is now serving a TLD, it becomes a sort of different matter than a lot of the other stuff we're talking about, sort of operational. So I wonder what that infrastructure looks like and just as a high-level table contest for the questions, taking whatever else. So I wondered, what sort of correlation has anyone done with when one of these registries is in a potential EBERO state where you might consider enacting EBERO with other observables? Like you see traffic patterns differ or people doing any kind of conscientious [moderating] to sort of see what does stability or instability look like for example at the root when we have a TLD that is in its [degraded] state so we know whether we should treat it like an emergency. Basically

there's a lot of stuff that we have to learn so I wonder if people are tracking that and looking.

And then finally, there's an interesting comment about can you successfully transact an EBERO on a situation if there are science subdomains? I'd be curious what the thinking around that is a little bit more. So those are my three sort of high-level questions. And I'm happy to [inaudible]. I don't want to use up all the oxygen but –

CHRISTINE WILLETT: So the first question was about the infrastructure of EBERO operations. So we have three contracted EBERO providers that are engaged to provide EBERO services. They are Nominet in the UK core and then EU and [CNIC] in Asia Pacific. And the requirements – I forget what was publicized about the EBERO operation. Francisco, you'd know more about the technical infrastructure of those three providers. But it's not technically ICANN hosting the TLD.

FRANCISCO ARIAS: Right. Those requirements were published in RFP. We did a public RFP. Must have been 2013 or 2012. I don't remember the specific date. So in that, we listed the requirements and we have three EBEROs as Christine said, and they are publicly set in the EBERO website.

You ask also about whether we were doing some analysis to perhaps infer what's going on or be prepared to identify these issues before – sorry.

ERIC OSTERWEIL: Yes, I just want to clarify that it was more sort of like, has anyone taken on responsibility to look at the different observation spaces we have and understand what is the impact of these degraded registries? We haven't enacted EBERO, but can we look and see how dire is it?

To James' point, he brought up – you mentioned there's a formalized process for like what's going to happen at what point and what decisions are made, understanding what the effects of the global ecosystem are as TLDs are in these various states.

It's sort of like we can look at the various measurements that have happened and understand how urgent this is. And I think I misstated my earlier question that your answer is really [inaudible] table until we do one at a time. So yes, I'm more interested in, are people looking at this like how important is it for us to go to EBERO when these registries are suffering? In other words, what's the impact to [inaudible] parties, for example?

FRANCISCO ARIAS: I think that's the primary reason for the existence of EBERO. We're trying to minimize the pain of the users. I'm not sure if that's where your question is going.

ERIC OSTERWEIL: Yes. Sorry, I feel like I'm dragging this into the weeds, so I'll stop if at some point it feels like we're getting too down and dirty. But if we say after N hours, it's an event. And after – to James's sort of strawman – two times N hours is a really big event. It's like, why is it N hours and

two N hours? In that time, if something really bad is happening, should it be like N/2, and should we absolutely at N hours be pulling the trigger? Because by the time it gets to two N hours, we've lost this many billions of dollars in ecommerce or something like that.

I'm just asking, have we thought about putting rigor behind some of this and making it more measurement-driven, for example? That's kind of where I was headed.

FRANCISCO ARIAS: So [inaudible] interests were defined during the new gTLD process, and I don't recall what was the reason behind those numbers. At the moment, the ICANN organization has not much of a say there. The emergency thresholds are what they are, and we have to execute to them.

The decision making that we have is whether at the point where we reach the emergency threshold, for example if there are no registrants, then again, what's the point? If there are registrants, then we have to decide, for example, what is less painful to the registrants and the end users. To keep it what it is knowing that or having some assurance from the current operator that service is going to be restored and say two more hours? Or initiate an EBERO transition that will take a day or two? That kind of thing.

ERIC OSTERWEIL: That's a great answer. I really appreciate that level of insight. That was kind of where I was headed. That's going to be an optimization problem,

who's bearing how much pain and at what point. That's a nice segue back to me clarifying my earlier question, which I think I just didn't say very clearly.

At some point, I'd be interested in hearing what sort of analysis or what sort of documentation or understanding is around who are the people and the processes and the systems that are involved when an EBERO event may or may not be happening, including what you mentioned, robo calling and outsource knock. Those are things that are relevant for business continuity or something like that that becomes an operational role of serving the identifier space.

And as much as you may not be doing 453 resolution, there's a set of processes that affects a 453 resolution. So at some point, I think we may want to know who the people – not by name, but are the processes part of a risk management framework, or where do they sit in this sort of infosec posture, and all of that. Can I spoof a robo call and trigger an EBERO event or something like that?

And sorry if it wasn't clear before. I'm sure it wasn't, but that's where I was headed.

FRANCISCO ARIAS: Yes. So I think your question is perhaps beyond what I can answer, but one thing I can answer you immediately, spoofing a robo call would not trigger an EBERO event.

BOBAN KRSIC: Hi, Francisco. Just one question, only to understand it better. You talk always about potential EBERO events. Have you ever tested end to end a new gTLD transition? With all entities. Because we talk about EBEROs. We have EBEROs, yes. We have DAs, and I found nothing in the history about such a testing of such a case. Can you clarify it?

FRANCISCO ARIAS: Yes. As a matter of fact, we just completed a tier, we call it EBERO exercise, a few days ago. So as of today, we have done one such exercise with each of the three EBERO providers that we have, and so we took advantage of the fact that there are new TLDs that wanted out of the contract. They wanted to terminate their contract with ICANN.

So in those three cases, we asked the registry operator if they were willing to let us play with their TLD, and they kindly agreed to that. So at the very end of the six-month window where that process that's indicated in the Registry Agreement.

So at the very end of the life of the TLD, we triggered – or we simulated as if there was a failure in DNS service for that given TLD, and we executive the full EBERO transition process. And that includes all the actors, as you said. For example, we simulated the worst-case scenario of an EBERO transitioning in which a registry is either unresponsive or unable to respond.

So we went to the [escrow agent] and asked them for the recent data, and with that we recovered the TLD functions. So we [inaudible] the full exercise. We went to IANA and asked them to change the [inaudible] in the root zone, etc. But we did the full cycle at [inaudible] times one

which of the providers, and I have to report that all of that was successful.

We have presented about this topic publicly. I remember doing a presentation after the first time we did this exercise. I think it was in Marrakech, and we were invited to present again in Abu Dhabi. So that's in the ccNSO Tech Day, and so there is going to be a presentation on the cumulative of the three EBERO exercises we have done.

UNIDENTIFIED MALE: Just one question. Do you have that presentation somewhere publicly available? Because we have [inaudible] meeting Tech Day. So we [couldn't] be present.

UNIDENTIFIED MALE: Really?

UNIDENTIFIED MALE: Yes.

FRANCISCO ARIAS: Sure. All the presentations from the ccNSO Tech Day are made public, and we can share that with you.

BOBAN KRSIC: You said your simulated it. What was the requirement from the time perspective? Because I read that you have 24 hours for a transition from

initiated, to collect the data from the DA to give the data, the EBERO and to make it public in the zone. And 24 hours are high requirement, and would it be also possible to not only simulate it, to say, "Okay, let's operate it and show that it really works in real life?"

FRANCISCO ARIAS: To be clear, what we simulated is the failure on the service.

BOBAN KRSIC: Okay. [inaudible] Everything else was for real.

FRANCISCO ARIAS: Yes.

BOBAN KRSIC: In 24 hours?

FRANCISCO ARIAS: I think you're referring to the SLA for the data escrow agent to release the data, right? That did not happen within the SLA. We have reported that publicly in the [inaudible] we did another presentation in Madrid in one of the events. That covered the second case. And so in both cases, the data escrow agent took a few more hours than their SLA. And as a result, we launched a small project to work with the data escrow agents to see how we could improve that situation.

BOBAN KRSIC: And what's the service level for the whole process, for a transition? Do you know?

FRANCISCO ARIAS: Sorry. It will take me a few minutes.

JAMES GANNON: I have a follow-up, and then a separate question. The follow-up is my understanding is you didn't simulate a scale however with regards to domains, because these were essentially empty TLDs.

UNIDENTIFIED MALE: [inaudible]

JAMES GANNON: So are you able to say publicly how many domains you simulated within that exercise and if you have plans in the future to simulate at a larger scale for the failure of a larger TLD? Because that goes to the ability of the EBERO operators to actually scale rather than just the process. And then I have a separate question after that.

FRANCISCO ARIAS: You're right, James. We used the data from data escrow agents, so in all these cases, they only had one domain name. nic.tld, obviously. And now in terms of the scale, we not only do the EBERO exercises with the EBERO providers, we also do annual inspections. That's the term we use. And they are a subset of the process of the EBERO transition.

For that, we use synthetic data. We create deposits, and for that we use more domain names. I don't honestly remember at this point how many we use, but it was more than one. So that's what we have done on that.

JAMES GANNON: Thanks. And I swear you've read my mind. I understand you do the annual inspection, but a question that I've asked for a number of years now is, the EBERO operators were selected back in the end of 2011, and ever since then, I understand that there is an inspection done, but it has never went into the security and stability of the actual EBERO operator. That is not part of your process of maintaining these three vendors as EBERO operators.

Can you speak to how we can have what I would see as a pretty critical part of Internet resiliency not be I suppose I would say audited as, "Okay, are they ISO certified? Do they have internal security processes?" That was part of the original RFP, definitely, but some of our EBERO operators – and I'm sure everybody knows what I'm talking about – have went through major crises of confidence in security world, yet there has never been an assessment on whether the EBERO operators we have should continue to be our EBERO operators, and is there a plan to review that situation, be it through a new RFP or be it through a more stringent review of the EBERO operators.

FRANCISCO ARIAS: We have multi-annual contracts with the EBERO providers, and if memory serves, two of them come – I'm looking at Christine, hopefully she would remember. I think two of them, their contracts end next year

I believe, and so that we need to address the process and potentially do an RFP to either renew with them or get new providers.

CHRISTINE WILLETT: Absolutely. I would also add that Francisco is more familiar with the registry operators, but these providers all also do act as registry providers, registry operators themselves for their own TLDs, so despite whatever crises of confidence in the past, they do have a track record and they're currently operating TLDs for themselves or for other customers. And as Francisco said, we do have procedures to monitor and exercise them on an annual basis. And I do believe that by the time we executed the agreements with them – I think it was 2013 – and so I think with five-year agreements, they would be coming up this fiscal year. So yes, we are in discussions, as we do with various vendors. We have multiple vendors that provide critical services to keep gTLD operators running, TMCH, TMDB, so we do look proactively at that.

One of the things on the table is potentially looking at resourcing for EBERO providers, making sure that the current requirements for EBERO providers are aligned with the current situation which to change, it's very different than it was – it could have been in anticipated in 2011 in terms of the scenarios of registry operations today. So we'll make sure that current requirements are met as we look to offer this service going ahead.

BOBAN KRSIC: Thanks, Christine.

JAMES GANNON: Sorry, just a really quick follow-up. This is an ancillary question. Has ICANN's process for engaging third party vendors for critical vendors changed since 2013, do we know? I.e., is it more stringent now than it may have been back then?

CHRISTINE WILLETT: I would say overall, we have more robust internal processes around procurement than I recall being in place five years ago, given the efforts – and we have a very robust procurement process, but the same requirements of expenditure thresholds, getting Board approvals, getting executive approvals. By and large, those are still in place as they were five years ago.

DENISE MICHEL: I have I think a higher-level question. And this has been really valuable to be able to delve into this area in particular. Thank you so much for spending time with us doing this.

We're obligated to assess the implementation of the first security review recommendations and their impact. In line with that, Recommendation 11 of the first security review says that ICANN should finalize and implement measures of success for the new gTLDs that expressly relate to the security, stability and resiliency-related program objective, including measurements for effectiveness of mechanisms, to mitigate domain name abuse.

We received sort of a staff paper that mentioned sort of an inventory of numerous activities that could be used and/or have the potential to support to be defined SSR objectives for the New gTLD Program. It appears that the New gTLD Program specifically doesn't have finalized SSR objectives that are then quantified and measured. I wonder if you can sort of elaborate on the [staff] activities and thinking and plans in that broader area.

NEGAR FARZINNIA: Yes. Thank you, Denise. If I'm not mistaken, that recommendation was one that we had actually hired a third party to help us evaluate the list of – you're referring to the recommendation for which we had a report with a list of possible areas that could be used for a system, but it wasn't finalized, and we had actually requested this review team to look at that list further and maybe make a better determination as to the appropriate items.

I'm not sure if the SMEs that are here are able to address that question, as we had, again, a third-party contractor help us out with that, but we're happy to take that question back and see what answer we may be able to provide to you on that.

DENISE MICHEL: Thank you. Christine, do you have other things to add about sort of SSR-related activities within your program? Any additional things that you're undertaking or plan to undertake related to this?

CHRISTINE WILLETT: Thanks, Denise. I apologize, I have not seen the staff paper that went to the SSR2 team and briefing against all of the SSR1 recommendations. As we've talked in the past, the New gTLD Program has implemented a variety of capabilities and functions to ensure stability, security and resiliency. Not just of course the EBERO program and the SLA monitoring system, but a number of efforts from technical evaluation, pre-delegation testing. Various aspects of the New gTLD Program are really focused on a route of ensuring the security, stability, resiliency of the DNS or the numbering system. But in terms of the metrics, I would rely on Negar and the team's perspective on recommendations there.

DENISE MICHEL: Thanks, and maybe asking a more specific follow-up question. I guess given the failure numbers and emergency failure numbers we've heard, in looking at the SSR-related reviews of the applicants and the interaction that staff has in helping registries meet their commitments, any additional thinking or plans in that area as you prepare for the launch of another new gTLD round?

CHRISTINE WILLETT: Thanks, Denise. One of the ideas that has been discussed to some degree with I believe the Subsequent Procedures PDP Working Group, also within the Registries Stakeholder Group, is the idea of a backend certification program, a registry service provider certification program. I think some of those discussions may still be ongoing. I'll let Francisco speak to the current state, if any, of those discussions.

But in concept, when that was being discussed, the concept of certifying these providers ahead of time could potentially offer the benefit of not just performing a technical evaluation up front, but performing a series of tests with them to ensure – to go directly through the SSR goals. But to also have that specific contractual relationship potentially between ICANN and those organizations so that we could work together with them as they serve registry operators and ultimately registrants.

So that would be another dimension in think that the idea was that that could augment our existing suite of offerings and capabilities that we put in place to try to drive SSR.

DENISE MICHEL: And just a quick follow-up. Setting aside what a policy development group decides to do or not do in terms of certification, I would think that staff has a fair amount of flexibility on how it sort of interacts day to day with new applicants, and to new registry backend operators. Are you also discussing at the staff level – understanding that this very first round of new gTLD applications has served as an enormous learning opportunity for everyone and in particular staff, and I imagine there are many things that you will be tweaking or changing or adding to when and if a second round launches, but it seems to me that staff has quite a bit of flexibility in the resources it brings to bear and how it interacts with the new applicants and the new registry backend service providers and operators. So setting aside the policy issue, do you look at the [staff fund] failures so far, and does that cause you to look at changes in resources and approaches you take with these new applicants? James? And then Christine.

JAMES GANNON: If I can build on exactly the same question. For example let's take the [11] emergency threshold breaches. Has there been any proactive – what I would call problem resolution to say, "Okay, here's something, a trend that we have noticed. It sounds like there probably [will be a] trend, and now we actively have staff reaching out to the whole group of registry service providers to say, 'Here's something we know is now an issue?'" It might be an issue in 10 or 12 out of the 30-odd, but "Here's something you need to be aware of. Here's something that you need to spend more time working on, because we're seeing this as a [inaudible] problem across the service providers."

FRANCISCO ARIAS: This is what I was referring before. We did raise this topic with the Registrant Stakeholder Group, and they have a discussion group. We're a part of that discussion group where we are trying to find to also improve this specific situation.

I think that was the question, whether we are doing something. We raised this topic with the registries. And I said, well, it's also being considered by the gTLD Subsequent Procedures. In that specific place, we provided our specific suggestions to the Work Track for it in response to their request on when will we recommend the process in new rounds of gTLDs, how the approval process, the relation process of these new gTLDs could be changed from what was done in the subsequent round. And we've made the recommendations on things like it would be may be a good idea to have some [virtual] process for

RSPs, changes to how PDP is done, how the technical evaluation is done, etc. All of that is public information and we can share that with you if it's of any interest.

DENISE MICHEL: We're over time. I guess we can always I guess put our questions in writing, but I think in part, what this raises for me is I think some additional questions. And I really want to set aside the policy development issue and really focus on what staff has the ability to do in implementation and just the normal interactions with registries.

I'll put that question in writing since we're out of time, but I'm also curious about the – we had a conversation with a particular registry about ideas about what can be done. You have ICANN the organization with responsibilities to the Internet's unique identifiers and SSR responsibilities broadly, which is a different set of concerns and responsibilities to an individual registry that may be much more focused on what, say, their bottom line is or other considerations in launching a specific business completely separate from the broad responsibilities of ICANN. So I guess what I'm looking for – or wondering if there is something additional that you're doing that addresses the broader stability and security of the Internet in this area.

CHRISTINE WILLETT: I'll take that, and then maybe I'll go back to one of the earlier questions. Just to be clear, there is a suite of functions we perform driving towards stability, security, resiliency. It starts with technical evaluation of registry operators. It goes through – if you're not all familiar – some

very rigorous pre-delegation testing efforts which Francisco fought vehemently with some of the registry operators to ensure that every TLD every time goes through pre-delegation testing and the [import there] that that testing remains.

On an ongoing basis, we have had – although we're not evaluating new registry operators at the moment – on occasion existing registries due to the new backend, or tomorrow a new service provider could crop up, and they are put through the same rigorous tests that any operator previously would have been two, three years ago in the past.

And then when registry operators notify us of a technical change to their operation, that individual technical change frequently going through our RSEP process is evaluated and could require a technical test and could require additional technical testing.

And then I think just one situation – I forget if it was James or you Denise suggested sort of a hypothetical situation if we became aware of a problem that was perhaps challenging. We noted it was challenging. One registry operator who we thought it might be pervasive to challenge other registry operators. Granting that it's a hypothetical situation, I would offer to Francisco, my guess is that you do your best to engage with the Registries Stakeholder Group and existing registry operators to convey as much information as possible about that hypothetical situation so they could route that off and get that issue addressed as quickly as possible.

FRANCISCO ARIAS: Yes. Of course.

JAMES GANNON: Yes. And just a quick follow-up. I'm going to ask a slightly loaded question.

DENISE MICHEL: [inaudible]

JAMES GANNON: I'm hearing a lot of when you want to do things, you're going through [inaudible] or you're going through PDP, the bigger, more broad community forms. Do you feel empowered to reach directly out to a group of registries that you feel have issues and try and work with them directly on the operations side rather than maybe the policy side?

There are two aspects to any of these problems. There is usually something that needs to be fixed in policy. It either needs to become a looser SLA or it needs to be something – the whole RISG needs to look at. But then there's also the ability – and I want to ask the question, do you feel you have the ability to reach out and work directly with individual registries to try and assist them where you have this known problem that needs to be addressed?

CHRISTINE WILLETT: Yes.

STEVE CONTE: Just from a CC perspective – because I know that that's honestly your [inaudible] and in relation to some of the training that OCTO and SSR does, we have by far more requests from the CC community than the G community to come provide training around DNSSEC or DNS operations. We do a three-prong – or we offer a three-prong course of initial DNS operations, advanced operations, and then the secure operations which is basically how to be a registry operator up through and including DNSSEC.

Our request for – and our offers of assistance – from my perspective in the organization – are much more well received within the CC community. I think mostly because of – from my perspective – from financial resources and human resources as well. There's a slightly more limited pool of resources that a CC can pull from, especially in an emerging economy, than a G who theoretically has resources and has gone through the G application process.

So there are still places and times where we're actively involved in helping the registry or a registry operator improve their infrastructure, improve their services. And it's not always falls within the G space as well. Eric, and then James.

JAMES GANNON: Just very quickly, just a question of organizational split then. Let's say a smaller gTLD wanted that assistance and education, would that fall under OCTO or under GDD? Or both?

STEVE CONTE: I'm going to speak, and you can correct me if I'm incorrect in the process, but if we as an organization get a request for assistance for training, it'll probably come through either direct because of personal relationships, through our Global Stakeholder Engagement Team which the majority of our requests come through, or maybe through the GDD Team.

Off the top of my head – and that's not authoritative – I don't know of any G that has every come to OCTO or SSR Team to specifically request training. That's not to say that hasn't happened, I just don't have that on top of my head. Did you want to add anything?

UNIDENTIFIED FEMALE: [inaudible]

STEVE CONTE: Sorry?

UNIDENTIFIED FEMALE: [inaudible]

STEVE CONTE: Absolutely. If there was a request, we would deliver training, we would work out how – everywhere that we go, we try to do the most bang for our buck, making sure that we're not just doing a one-on-one training with a specific registry. We want to include as many other players in that session as possible. Eric?

ERIC OSTERWEIL: This just gave me a follow-up question. I hope it's short, and we can punt it for later if it's not. Taking the situation and reversing it, we mentioned RCAs a minute ago. I think Francisco you mentioned RCAs were requested from – in the case of EBERO. And [inaudible] can be really sensitive, so certainly if they were made public, that would be awesome, but probably is a lot to ask for. Is some level of digest being created for the common faults, the common problems that have caused outages or problems that can then be digested as sort of like read-ahead material as opposed to requesting training? Those who are thinking of walking down this path, these are the known things that we've seen people stumble over at the [various] rates and how to avoid them proactively. Something like that. Heads-up doc.

FRANCISCO ARIAS: Yes. And it has been done. We have had sessions. I can't remember when this was, probably – I'm sorry, I just don't remember. But we have had sessions in which we discuss at a high level the kind of issues we have seen with the registries. This is what the root causes that we have seen reported and the issues we have seen related to the SLA Monitoring System. And so you may want to take a look at these kind of things so that you [inaudible]

ERIC OSTERWEIL: Yes, that's great. If someone has a pointer to that. If it's just me who doesn't know about it, whatever, but I think we'd love to get a look at that and see what we can do with it. Thanks.

CHRISTINE WILLETT: Is that one of the registry operator roundtables? One of those sessions at an ICANN meeting?

FRANCISCO ARIAS: I don't remember which session that was.

CHRISTINE WILLETT: We'll find the sessions. I don't know if they were recorded or not, if they were public sessions. We'll go back and find pointers to that, but I can recall Francisco hosting that on more than one occasion.

JAMES GANNON: Just very briefly, have you ever thought about moving that from, "We present at an ICANN meeting" to an, "Okay, once we identify a consistent problem," we do up a proper root cause analysis and potential mitigation actions, put that somewhere on the website where you can say, "Okay, here are the recurring problems that we see?" And more kind of a structure that as these things come up, rather than waiting for maybe the next ICANN meeting [inaudible] presenting on it, more of an, "Okay, here's our current info on what the challenges potentially for RSPs are, and our suggested directions to go with to solve them."

Of course, that could not be just for emergency threshold breach, it would also – the smaller SLA breaches. Have you ever thought about a

process around reporting and guiding the RSP community on those issues?

FRANCISCO ARIAS: No. Thank you. That's a very good suggestion. We have not thought of that.

STEVE CONTE: Pin that in as a possible recommendation then. Just to clarify your question, the intention is to not name and shame, but to do a lessons learned and how to continue and to continue to run the system in a more robust manner? Okay.

ERIC OSTERWEIL: Yes. That's absolutely – actually, even more so, it's like you can start to understand what it takes to run a registry if you look at where things have fallen down before any kind of [inaudible]. It's not a name and shame, this is a proactive measure that you can do now with real data that later on could factor into something like assessing qualifications. So yes, this is basically a step one of what I think might be a really good set of things down the road.

BOBAN KRSIC: And maybe it can also end in an adjustment of the threshold, not to talk about potential EBEROs then cases, then define a threshold when we are sure that it's an EBERO case in place.

So for me, it's hard if I read that [inaudible] the potential EBERO cases and we don't have – not relevant. Maybe we can adjust this threshold to another level and say, "Okay, when we reach them, then we have really one case." And then we have to move along our procedure for an EBERO case.

FRANCISCO ARIAS: I'm sorry, I'm not sure I understand the suggestion.

BOBAN KRSIC: An adjustment of the thresholds. Because we are talking – no? James.

JAMES GANNON: My understanding is because it's community-developed policy, you can't change those thresholds, correct?

FRANCISCO ARIAS: Oh, okay. Yes, if we're talking about changing the emergency thresholds or the service level requirements, those are baked into the Registry Agreement. So the ICANN organization cannot change that [inaudible]

ERIC OSTERWEIL: Yes, so just to stand right in the middle of all this, we could do measurements to reflect what the over and under is on breaches at certain points in order to say back to the community, "It looks like at this point, you've now crossed a threshold where it's costing more than you're gaining by not going forward. Therefore, we recommend etc."

So I think by sitting with the current thresholds that we have, we have the opportunity to do measurements that might inform a sort of a measured reason to adjust in the future.

ZARKO KECIC: Yes. How you are assuring that SLA are actually emergency threshold is reached from a technical point of view?

FRANCISCO ARIAS: I'm sorry, what was the question?

ZARKO KECIC: Are you 100% sure that the emergency threshold is reached because of actually – from how many points are you doing measurements and monitoring?

FRANCISCO ARIAS: Yes. In all the cases that reached the emergency thresholds, the [inaudible] denied there was an issue, and we worked with them constructively. So I don't think there was any doubt that the issue was happening.

To your technical question, the planned system that we have has approximately 40 probe nodes. I say approximately because these systems, sometimes one probe node goes down and another comes up. So it's around 40 probe nodes that are distributed throughout the Internet. According to the SLA, we have the requirement to put the

probe nodes in the case of the DNS where the most Internet users are. There is guidance in the SLA where to place them. The [probe nodes] we're trying to follow that as much as we can, given the availability of where to put these servers. So we have our own 40 that are measuring at any given time each of the TLDs in the root zone.

UNIDENTIFIED MALE: Negar.

NEGAR FARZINNIA: Thank you. I just wanted to bring the agenda back to everyone's attention. We have Brian here only until 2:45. I think he's actually staying past that point. If you don't mind, if we can cover the statistical analysis of DNS abuse study while we have him here, and after the break, we can continue the rest of the conversation as we have the other SMEs remaining with us until 4:00.

That's one point, and then the second one is, some of the questions that were discussed here today, I believe they were asked during the Johannesburg meeting and we have provided answers to them in writing. I'm going to ask Yvette to post the link to those questions and answers for everyone's reference, just if you need further details on the answers that were already provided.

Thank you. Yvette, if you could please post the link to those questions and answers, I would appreciate it.

JAMES GANNON: My last question is a yes or no, and you may tell me that you don't want to answer. Do you feel that the emergency thresholds are actually fit for purpose and appropriate, or should the community look at revisiting them?

STEVE CONTE: Is that a question that ICANN as an organization should be answering, or should we be providing data in which the Empowered Community should be questioning that themselves?

DENISE MICHEL: Can I offer an answer? I think as experts in this field, we'd be interested in what the unofficial opinion is of ICANN organization staff as I think one of many factors that go into an ultimate decision that would involve the community and I think other players as well.

STEVE CONTE: Then before I offer you the mic, with the understanding that it is at this point only opinion, then I will get out of the way between Francisco and the mic.

FRANCISCO ARIAS: I don't think I have done enough analysis to answer the question. I don't want to just offer an answer without analysis. We have some data and I'm sure if someone – [there's some] analysis group come up with an answer to that. I just don't have it in me yet.

CHRISTINE WILLETT: May I? I think it's a really interesting question, and I think it goes beyond just what are technically appropriate thresholds balanced with the impact to registrants and end users. How do you formulate a multidimensional set of thresholds that the downtime impacted users potentially, and factoring in that cutover time. And at some point, as Francisco said earlier, when it's clear that the cutover is going to take longer than the time – again, if we have confidence in the registry operator – that they're going to resolve, the ultimate impact to registrants and end users ends up being in the balance, and there is a judgment. So I think it's an interesting question, and it doesn't just go to what's technically appropriate, and how do we hold a firm bar? It's also the practical impact of implementing an alternative solution.

UNIDENTIFIED MALE: [inaudible]

BRIAN AITCHISON: Okay. All right. My name is Brian Aitchison. I think I've met you all before briefly, and you've seen me on calls before. But I've helped manage this statistical analysis of DNS abuse in gTLD study that I know you've all heard about. So my presentation is actually very quick. I know you guys are constrained for time, and I was asked to just kind of give some high-level bullet points.

You can see a more comprehensive presentation [while] we gave a few webinars to the community a couple of weeks ago. We'll post the link in

the chat and/or share with the mailing list as you like. It's about 45-minute presentations and they'll go into much more detail than I will here. Also of course read the report if you haven't already. That's really the best thing to do rather than listen to my five-minute spiel here.

But I'll give a little bit of background. I won't go too much into the methodology because I think you're probably most concerned with the results of the study, so we'll hit on that.

Essentially, to start off it was requested by the CCT Review Team, and they needed a set of descriptive statistics to measure abuse in the domain name system. Part of their mandate was to examine malicious abuse issues in the New gTLD Program and to analyze the effectiveness of safeguards that were part of the New gTLD Program.

So to do that, they needed to see if safeguards presumably would have an effect on abuse rates, so they needed to see what those abuse rates looked like. And they also saw this as a proxy for trust, which is in their name, so if you see a high abuse rate in a certain TLD, probably can infer that that's not going to be a highly trusted TLD in some way.

So what did we look at? We looked at phishing, spam and malware in new and legacy gTLDs from 2014 to 2016. And methodologically, the very short description is that we cross-referenced WHOIS, zone file and historical blacklist data to kind of filter out what the sort of bad guy domains were.

We also employed an inferential statistical analysis, kind of a driver analysis to look at the effects of DNSSEC, domain parking, registration restrictions and TLD sizes on abuse rates. We had [inaudible] the Delphi

University of Technology help conduct the study, and the CCT RT will look at the report, review it and make recommendations based on that. They're currently in that process right now.

Now the good stuff, the findings. Again, very high-level findings here, and this is the broad finding, is that there are high levels of abuse, but it tends to be concentrated to a relatively small number of new gTLDs.

The abuse counts, which is an absolute number of abuse domains, are relatively constant, and overall higher in legacy gTLDs. This is to be expected, of course. There are many more registrations in legacy TLDs. What we do see in terms of absolute counts is an upward trend of abuse in new gTLDs. Again, no surprise there. They entered the picture in 2014. We could expect there's some amount of abuse that's going to be going up in new gTLDs.

Now when you look at rates, which is based on the abuse domains per 10,000 ratio, rates of phishing and malware tend to be lower in new gTLDs than in legacy gTLDs. There are exceptions and spikes based on TLD and timeframe, but the rates appear to be converging to similar levels by the end of 2016.

One of the most interesting findings I suppose is that spam counts – and remember, counts is an absolute, not a proportional number – are actually higher in new gTLDs compared to legacy gTLDs. So when you balance that out with the number of registrations or compare that with the number of registrations in legacy, it's quite surprising to see that far more spam is actually coming from new gTLDs.

A few of the other more minor findings, privacy and proxy services don't actually appear to be associated with any abnormally high levels of abuse. If you look at registrar location, the U.S. and China tend to be associated with the highest absolute amounts of abuse. That's not surprising given these are the sort of largest markets for registrars.

But when you look at rates, again abuse domains per 10,000, Gibraltar comes on to the picture and there's one particular registrar that's headquartered in Gibraltar, so it kind of – this geographic analysis has helped us kind of pinpoint where some of the abuse issues may be stemming from.

The researchers were able to distinguish between compromised and maliciously registered domains. This is sort of – they employed a few general heuristics to help them do this, but essentially the way they were able to do it is if a domain was flagged on a blacklist within three months of registration, if it had a misspelled variation of a brand name or it contained a brand name, they can infer that the domain was most likely – most likely being the keyword – maliciously registered. Again, this is sort of a broad exercise to help characterize these kinds of domains.

So when you look at compromised domains, it's much higher in legacy gTLDs, and maliciously registered domains tend to be much higher in new gTLDs. And another final important finding is this finding from our inferential analysis was that registration restrictions appear to have the strongest statistical effect on reducing an abuse rate. So the more restrictions you have in place, the less abuse you're going to experience, and vice versa.

So those are the high-level findings, and I'm happy to answer any questions on the study. But otherwise, I would encourage you to read it and listen to our webinar.

DENISE MICHEL: Thanks. So I think my first question is, why didn't ICANN staff conduct this study after a critical mass of new gTLDs were out there and underway? Why is the study – this is a genuine question – coming from an external community Review Team contract with a consultant? It seems to me – so we're trying to work our way through the SSR1 recommendations. In part, that includes requests for SSR objectives in the New gTLD Programs, and metrics. So I don't know if this thread seems logical to you, but in my mind, it seems potentially one of those SSR metrics involving the New gTLD Program would be abuse in new gTLDs. So I think my first question is, was it ever contemplated that staff just take on the responsibility of doing regular statistic gathering and analysis of abuse rates in new gTLDs? I think that's a two-part question. And if not – so I'd like to understand the thinking around that, and then if it didn't occur to staff, I guess why not, and now that this abuse report, the first one has been issued, is staff rethinking the need to collect this type of data and issue this type of report on an ongoing basis? Thank you.

BRIAN AITCHISON: I may have to punt a lot of that to someone else, but I can only speak in the context of the CCT Review Team that as far as I know, there was no community mandate to conduct these kinds of studies. But again, don't

take that as authoritative. I entered the picture personally about two or three years ago, and that's when the CCT kind of came up with this recommendation that we carried out.

DENISE MICHEL: I'm sorry, just to clarify, so are you saying that staff doesn't have the authority or ability to collect data in this area?

BRIAN AITCHISON: No, I'm not saying that. But Steve, I think –

STEVE CONTE: So we've talked about DAAR already today and we had it in Madrid. There's clearly a need or collecting and analyzing data, and whatever barriers that John Crain and Dave Piscitello had to get through to get DAAR going, they got through. It took longer than people wanted, and probably not long enough for other people who might be bad actors.

So I think to answer your question of, "Can we?" I think the answer is yes, we can, with the caveat though that even DAAR is using external data. So unless there's a significant effort – I don't even want to bring budget into it – to say ICANN should be the one who is out there and have nodes all around the world to monitor specific data types, then we as an organization will be reliant on external parties, either be it to do a complete study, or to at least get the feeds for the DAAR project in which then we can use that data to evaluate abuse and various things.

DENISE MICHEL: Just to clarify, are you saying that staff does have the authority to collect this data, or contract people to collect this data unilaterally, or not?

STEVE CONTE: We've been given the go ahead to collect – this data is – and I'm not sure what you mean by that. We have the authority and the go-ahead to collect data through these data feeds in DAAR, the 13 feeds or whatever it is, in order to analyze and begin to analyze abuse and trending and things like that.

Whether or not – it's not [inaudible] your question of why we didn't do it in the past. I think it was in some way a different world in the past, and we've got a different environment now and people are more data-oriented now than ever before. Eleeza, do you have your hand up? Do you want to add to that? Yes.

ELEEZA AGOPIAN: Excellent. This is Eleeza Agopian, by the way. Hi. I just wanted to add that this particular study was undertaken when we asked the CCT Review Team, so Brian and myself have been providing research support to the CCT Review Team. And as the Safeguards and Trust Subteam of that Review Team was looking at the question of trust and how do you measure trust, which I think you'll all agree is a difficult thing to pinpoint. One of the areas they really wanted to dive into was the impact of the safeguards themselves, and see if we could correlate that to trust in some way. So this was their heuristic for measuring that, which is where the study came from, to give you a little more context.

DENISE MICHEL: Thanks. Now I understand the context for it. I think the context – I'm taking that abuse report and putting it in the context of SSR and the New gTLD Program, raising the question of – it makes sense that abuse metrics would be one sort of metric for SSR measurements in the New gTLD Program. I think that's the context I'm putting this in.

STEVE CONTE: I'm not sure how to respond, because at this moment we are now – so we can talk about the past and we can dig to why we didn't or why it took us so long, but we're there at some measure, and if we're not in the SSR1 Implementation Subteam – which that may or may not be a valid question – moving into the ICANN SSR Subteam, the fact is that we are doing that now. And so I'm not sure I'm understanding your line of questioning, of looking at the past when we now have a subset or a baseline [where we can tweak] and start moving forward and start collecting data.

DENISE MICHEL: Sure. I understand your point. I guess let me take a different approach. Although this particular abuse report did not make any conclusions regarding the effectiveness of safeguards that were introduced by the New gTLD Program, I guess one would come to the conclusion after reading this report that if the new gTLD safeguards have been effective, there would have been an observable reduction in the level of malicious registrations in new gTLDs compared to the legacy TLDs. That of course wasn't the case. Is staff drawing any conclusions, or has the result of

this study generated any new areas that staff is exploring in terms of abuse and new gTLD safeguards? Anyone?

ELEEZA AGOPIAN: I don't think I can answer your question, but to your earlier point on ongoing – sort of coming back to this again and reviewing the numbers again, we do have a recommendation from the CCT Review Team to continue with these studies, to do so in coordination with the DAAR project. So kind of this is the baseline, and seeing where it might take us in the future. In terms of how that might impact things operationally, I'm not capable of answering this question.

DENISE MICHEL: This is an operational question, not a CCT [review] question. Thanks.

NORM RITCHIE: I [saw Denise's] team, their presentation just last week in Toronto actually on this report. One of the conclusions they did come to is that the rates were the same rate now between the legacy and the new gTLDs. So I think it'd be worthwhile for us to have a really good look at that report. It's very detailed, there's a lot in it. And there's a lot of depth there that we could look into.

STEVE CONTE: So Brian, just to give you [surprises] since we're live and being recorded, I brought that report up earlier today and suggested that we try to find

a slot in the plenary for SSR2 to have a more in-depth briefing for the Review Team if that's –

BRIAN AITCHISON: Yes. I have a presentation for it. I'm happy to give it anytime, but I figured we're short on time today.

STEVE CONTE: Denise, I'm not sure we fully captured the answer to your question, so I think what we'd like to do at this point then is capture it, make sure that we were expressing it correctly. We'll capture it out of the recording, make sure with you that we're expressing it correctly, and then bring it to people who aren't in this room and see if we can get an answer that is close to what you're trying to achieve.

DENISE MICHEL: Thanks. I appreciate it.

NOORUL AMEEN: Actually, the kind of incident you considered were spamming, phishing and malware, and these are under the classification of DNS abuse. If you considered DNS abuse, it would be pharming than phishing. Phishing is the kind of domain abuse, not a DNS abuse directly.

Second thing, the Internet spamming and phishing [inaudible] are very much related. So most of the spammers they will send phishing URLs to trick the users.

And third thing is that instead of all these things, there will be more incidents, whether DNS abuse is actually happening. And the kind of feeds you have taken from [inaudible] Spamhaus, APWG, these are good, but you could consider because browsers, Mozilla and Internet Explorer, they have a function to report phishing URLs by the users.

So there were a couple of [feeds] like PhishTank user browsers. We could consider that also for getting a better visibility of this study, I think.

BRIAN AITCHISON: Thank you. The parameters were set by the CCT Review Team on what activities we actually looked at. There was also a very large practical obstacle on the kind of historical data we could get. This typically isn't stored. Spamhaus and SURBL and the like don't typically store abuse data over the course of years, because that's huge amounts of data that they'd have to pay to store. So we had to get creative with how we kind of dug around for the data.

But yes, there's lots of potential for an expanded scope of the study, but this is I guess what we're able to do with the time and resources we're allotted, which was a year time to do a very large study of DNS abuse. So happy to take that kind of input, and perhaps do a more refined study in the future.

ERIC OSTERWEIL: Just real quick because I know we're over. And I definitely look forward to the detailed review. It looks like really good work and I think it's great

that you guys took it on and you did such a conscientious job and reached out for collaborators and stuff. That's great.

One of the things I would just sort of throw out there – and like I said, I look forward to the detailed review, so I don't want to get into methodology since we didn't talk about it. Being careful what your data sources are, going back and constructing a longitudinal dataset exposes you to potential errors that then become systemic, and then it becomes hard to control for them later.

So certainly going forward, like you said you want to go forward with more of an ongoing sort of thing, it'd just be interesting to cull the data sources a little more carefully, build your own sort of whatever. Because certainly if someone reports abuse anecdotally a long time ago and that winds up being sort of a control point for the trajectory you draw on your data, then all of a sudden your results are complicated.

I'm sure you guys looked at that. I honestly have had the report printed out for a long time and haven't read it, so I'm definitely not trying to criticize, but I think it's really great that you guys did such an in-depth, detailed study.

STEVE CONTE: Thank you for that. And we're open to the criticisms and the constructive criticisms, because we do want to continue doing this kind of thing. We are aware of the limitations of domain blacklists in particular, and we don't necessarily view this as a perfectly accurate view of abuse, but it's reliable.

And in terms of the data sources we were able to use, even if you take into account certain percentage of false positives or false negatives on those lists, you can still see general trends. And especially when you corroborate the lists against each other, one of the interesting things was that the findings – the trends are generally the same even from different lists that don't necessarily have overlap in terms of the lists are flagging.

So that's interesting in and of itself. An interesting way I've heard these lists characterized is perhaps there are false positives or false negatives, but these are the way that end users view TLDs. Whether it's actually spam or phishing methods that they're flagging, that could be wrong, but this is sort of the perception of the TLDs, that there's a lot of abuse coming from them. So that's kind of the way I tend to look at it at least, given the known limitations of blacklist data.

DENISE MICHEL: Just a quick follow-up on that. Why was it decided to do a sample of the gTLD data rather than scanning the entire – the legacy zone?

BRIAN AITCHISON: We didn't do sampling per se. we used a sample of new gTLDs. I believe it was 18 out of the 22. There was no data for .gov, .edu, .mil, and of course .arpa. Of the legacies. I said new, I'm sorry.

We had data for I believe it was 1196 of the new gTLDs that were available at the time, so I don't know if you're referring to statistical

sampling. We essentially had the entire population of data, so there wasn't really a sampling technique to employ.

The "sample" that we used was essentially most TLDs excluding ccTLDs that have been delegated at the time. The only reason a TLD wasn't included is if we didn't have the data for it.

DENISE MICHEL: Used the whole zone for new gTLDs and used sampling for legacy gTLDs?

BRIAN AITCHISON: I hesitate to say sampling, because that implies a technical, statistical technique. But we used –

DENISE MICHEL: [inaudible]

BRIAN AITCHISON: Right. As far as I recall, the data for the four TLDs that I mentioned was not available.

DENISE MICHEL: [inaudible]

UNIDENTIFIED MALE: We had data for all of they, yes.

DENISE MICHEL: [inaudible]

BRIAN AITCHISON: No.

DENISE MICHEL: [inaudible]

BRIAN AITCHISON: Correct. Sorry for the misunderstanding.

BOBAN KRSIC: So, time for break, yes? I would like to suggest – okay, we have a 15-minute break, and then we can come together with another one hour of GDD stuff and go back to work plan or the topics which are not discussed today.

I would say thank you to Brian and [inaudible] and see you in 15 minutes then here.

[BREAK]

ERIC OSTERWEIL: Mo, could I get you to unpause? So, we're back to it and I'm going to quickly pass this over to Boban. But just a little level-setting here is my understanding.

Christine and Francisco, thank you for staying and I'm hoping that this next session, we utilize you as SMEs to help the discussions for the work plan and stuff like that. Correct me. Unless you feel there is more need for presentation-like discussions, I think now is going to be more of a collaborative model. But with that, I'll turn it over to Boban.

NEGAR FARZINNIA: Sorry, I just wanted to make sure the recording is starting before we continue. Perfect, thank you.

BOBAN KRSIC: Yeah, welcome back. Thank you also from my side. It helps, really, to clarify some issues we had and I'll talk to James and I will give James the opportunity to introduce and this was [in the next slot]. To use the next 25 or 50 minutes to talk about some points which we have with these related topics here. So, James?

JAMES GANNON: Thanks. So there are a number of operational systems that support you guys in particularly the new gTLD space. There was the application system, there is a number of the GDD portal, and some of the systems behind this. Particularly during the new gTLD application, there were a number of security breaches or [inaudible] in those systems that had security implications.

Without going into the specifics of what happened, going from a set of systems that had security failures, let's call them, to where we are today, is there a process that you're aware of – and this may be split between you guys as the business SMEs and the IT Teams – on how, number one, those security issues were remediated, and then also how those systems are managed and operated from a security point of view now?

So we can start off with the ones that had issues and move on to what are the other systems that you guys manage that are used in more day-to-day operations now that we've moved past the application stages.

CHRISTINE WILLETT: Without talking specifics, and I don't want to speak for [Ashlyn] or Gary or anyone else from IT. I think they're going to be with you tomorrow.

STEVE CONTE: Yes, and just as a reminder, too, that we are open for session right now, too, so talking specifics, I think your question was more about generality, that there were issues in compromises in the past, what have we learned, how are we moving forward? Okay.

CHRISTINE WILLETT: So based on the incidents in the past, I can say that there is an even greater, heightened sensitivity to not just the security of our systems and user access, but from all dimensions, there's an awareness of the need to ensure that our systems are very robust, very secure.

We use a multitude of systems to get our everyday job done, not just the SLA Monitoring System that we've talked about at length today, the Centralized Zone Data System. We have our own operational tool based on the Salesforce platform.

IT has taken on a wide-ranging project that I think they'd be in a better position to address with you tomorrow about some of their activities. But I'm aware that they have implemented a variety of testing around user access procedures on existing systems. They've gone through reviews of existing systems. They go through extensive reviews on new systems, ensuring that they are restricted in the way that they should be.

So I think that the specifics, we don't implement. We just use. We're users of the systems. So I think most of those questions would be better addressed by IT and ITF tomorrow.

JAMES GANNON: So that's great and we will obviously address that tomorrow. As of more personal interest, is GDB [ops] involved in what I would call user requirement gathering for these systems because my understanding is most of them are pretty custom to your group, and how do you manage from the business ownership side of those systems, the evolution and the management of those systems? Are you involved in that on a day-to-day basis, ensuring that new requirements are gathered and implemented as to your needs?

CHRISTINE WILLETT: Thank you. Yes, for some systems to a greater or lesser degree. In general, as the business users of the system, we do articulate our requirements. We have a product management function that is responsible for gathering those requirements and translating them into technical details that our IT and engineering groups can implement. On occasion, we do have folks such as Francisco or others who, essentially, fulfill that product management role and may deliver those requirements themselves directly. But yes, we certainly have a hand in providing requirements as well as a level of user acceptance testing of systems and tools.

JAMES GANNON: And then my last follow-up, sorry, is there, then, as a result of that process, a formalized process whereby somebody on the business side, so separate to IT, gains ownership of that system and owns its development, so what I would call a business owner for an IT system? Or is that role looped into the product management role?

CHRISTINE WILLETT: I think it can be different situations for different tools. Sometimes it's owned by product management. Sometimes there's greater ownership by an operational team. We've had both scenarios, and over time, it also depends on the individual resources and who makes sense to be that subject matter expert or business owner.

STEVE CONTE: Just from a business perspective, regardless of services run and who the owner of the service is, IT still holds the machines and the budget for purchasing those machines. A budget is a budget. There's ways to move around things like that. But the owner of the servers is IT. The owner of the service might be Dave, the business unit – in this case, GBB – or it might be a combination of those two, or it might be a shared ownership between more than one business unit as well.

So depending on where that question lies, the ultimate people who service the machine are going to be IT and within the trusted network administration staff.

JAMES GANNON: Just to give you some context, for example, I'll give you a scenario where business ownership actually has an SSR impact. So for example, on some of the GD operational systems, it's very important from a governance point of view for the RTO and the RPO, so the recovery point and the recovery time objectives within a disaster recovery scenario, to be set by the business because you guys are the ones who know what your requirements are and in a system where IT has full ownership, they may not be in a situation to define those correctly and good IT governance says that decisions such as those, and again, poise or requirements, it's important that the business has not necessarily ownership, but governance over those requirements, particularly in a disaster recover scenario so that the IT Teams are implementing those to business needs rather than to what IT thinks you should need.

FRANCISCO ARIAS: I can speak for the systems that my team has since those are the ones that I know, so for those, we provide that function. We are the business owners and we also play a role in the interface with IT. Like other systems, we've very [inaudible]. We are more involved in that setting and specifically on the RTO and RPO, we define those to IT. We say, "This is what we want," and of course, there are some negotiations sometimes to get to something that is mutually agreeable, but we are the ones that request what is expected from those systems.

NORM RITCHIE: I have one question that I was picking up during our break. And looking at some BEROs and EBEROs, is it a requirement for an EBERO to also be a Back-end Registry Operator with current gTLDs, so therefore, you also have SLA methods on them?

FRANCISCO ARIAS: No, I don't recall any such requirement. At the beginning, we were thinking should we ask the opposite, that they are not a backend operator for a gTLD. But at the end, we decided that we could solve that problem by having multiple operators, thinking of what is the [inaudible] party is the party that you are trying to [inaudible].

NORM RITCHIE: Yeah, the reason that popped into my brain is that I think it was James asked a question about what tests are done for security on the EBERO operators. I thought, well, if they're actually backend operators and you

have the SLA, you kind of have a pretty good indication right out of the gate.

CHRISTINE WILLETT: Back in 2011, when the RSI/RFP was published, all the way to 2013 when we contracted them, we didn't have any new G registry operators. So yes, at the time, we didn't envision that they necessarily had to be, but it turns out that they are. So that's the current situation, but don't know what that might look like as we look ahead for the program.

NORM RITCHIE: Okay. Also, just in the realm of random questions, for the failures, for the registries that have failed, there's different reasons why they might have failed. It could be technical, it could be policy, it could be whatever. But it could also just be business, like they were not viable. Do you think that was the case in all cases where the registry was just not viable no matter what?

FRANCISCO ARIAS: Are you asking for the cases we have seen [case-to-case]? So those were [times] of failures. Something had gone wrong in their day-to-day operation. That's what happened.

NORM RITCHIE: No, but it could not be remedied? So, it failed to transfer, you invoke the EBERO process, and they're transferred away? Or am I missing something?

FRANCISCO ARIAS: So we have not had any [inaudible] transition yet. We have had 52 cases that reached the emergency threshold in which we decided that it was better not to do the transition officially because the [inaudible] group fixed the issue after [then we could] transition to a new provider.

NORM RITCHIE: Right. Okay, so I understand that. But given that we're into a renewal period for some of these registries now, so that's going to be a very difficult task for some of them, do you expect that there may be some people who will walk away?

CHRISTINE WILLETT: So we've already had 15 terminations, in that range, terminations of agreements. I can get the exact count. The exercise on .chloe was the last in that era. So we do have a number of registry operators walking away from TLDs, so we imagine it's no longer a priority for the business. Perhaps it's not viable or there's been some strange strategy.

NORM RITCHIE: Okay, one last question on that then. So for the next round, if there is a next round, I guess the point I'm getting at here is that there's going to be registries that fail not because of technical reasons or policy reasons

or anything else. They fail because the registry is not viable or is a bad business decision. Part of the application is actually assessing the business aspects of it, so do you think any changes are required in that or is that an out of scope question?

FRANCISCO ARIAS: I can only speak for the technical side. I don't know the financial review. I was not part of that. On the technical side, we provided recommendations to the PDP per their request. What I can say on business failures, at least I don't know what made the [inaudible] to decide to terminate their agreements. It may be they decided the business wasn't viable. I don't know what happened there.

And there's also other cases that sold their business to someone else. We saw those cases as assignments, and again, we don't know the reason why they did that because what I'm trying to say is that if there is a business issue, there are other options. You don't necessarily have to go to something that causes an emergency. You can usually know that you are not going to a good place and then have options to sell.

JAMES GANNON: A follow-up and then a second question. Thanks for [inaudible]. I have lots of questions on this.

So first of all, scarily enough, 27 terminations to date. But all of them have been .brand, which means they were essentially unused so we've never encountered a situation where a TLD under scaled use has gone through a termination, so it's still an unknown, unknown.

And then secondly, I forgot to ask this earlier, so under the EBERO contract, there is the annual inspection, which you guys are doing. There is also a clause in the contract for up to semi-annual third party audit of the EBERO vendors. Has that audit clause ever been exercised or is that something that is there for an unknown event to be audited, or what's the history behind that clause?

CHRISTINE WILLETT: So the history behind the clause, the intention was if ICANN had concerns about the operation of the EBERO provider and wanted to get more insight into their operations, perhaps after the annual exercise or for some other reason, then we could invoke this third party audit. But I can say, I don't know that we've said this publicly, but there haven't been such an audit of the EBERO providers to date.

JAMES GANNON: I'm going to ask another tough question. So with one of the EBERO providers, there was a major event that took place that caused the provider, as an organization, to lose a lot of trust in the community from a security point of view. Did ICANN consider exercising that audit clause with that specific event?

CHRISTINE WILLETT: I don't think I can comment on any specific situation.

NOORUL AMEEN: Supposing the vendor, the distributor [inaudible] to the registry providers, what kind of mitigation, because I have seen a document that experience mitigating [distributor in a lot of ways] as part of an ICANN document. So what kind of support do you [inaudible]?

FRANCISCO ARIAS: Sorry. Is the question, what kind of support ICANN would provide to a registry that is suffering the [inaudible] effect?

I don't recall that being specifically [inaudible] in the registry e-mail or some other requirement for ICANN to do that. Now when there is a specific issue that is going on, it is one of the questions that we ask to the registry if there is any [inaudible] that could happen or that they are willing to share, if there is an attack. But we are not set up to help with that.

NOORUL AMEEN: Because in an emergency [inaudible], the registry operator request for information [you lose] mitigation that was mentioned. That's why I asked the question.

FRANCISCO ARIAS: Oh, okay. I think you're asking about the requirement for the EBERO provider. Okay, sorry. Okay, right. So yes, the EBERO providers like the backend registry operators for gTLDs, they are required to have in place protections against U.S. and I don't think there are specific requirements on what those protections should be, but it's more about them disclosing that to us so that we know what they have in place.

NOORUL AMEEN: Is there any specific mechanics of that? These are the mitigation of [inaudible]. Those kind of [inaudible].

FRANCISCO ARIAS: No, I don't think they would be publicly available. They share them with us, but we don't make them public.

UNIDENTIFIED MALE: I had a question in the previous session, but I didn't have a chance to ask. How much of your data are used for compliance purposes?

FRANCISCO ARIAS: Are you asking about the SLAM data, the SLA Monitoring System?

This is basically a compliance system. I'm not in Compliance. I'm part of the GDD Team. So the way it works is both teams get the data, the GDD Team and the Compliance Team, and there's an analogy there to be made that we are playing the good cop and they are playing the bad cop. We are trying to help them to fix the issue while they [present] to comply with the agreement and so there is that sort of dynamic where we provide all the information we can to what we have seen in the issue so that they can fix the issue as soon as possible.

But our Compliance Team, they are independent and they do their function and they look for compliance with the agreement.

UNIDENTIFIED MALE: That was my first question. The second question, because I didn't get it probably I missed something when James asked you, what happened when they reach threshold in your measurement. Who makes the decision are you going to move registry operations to EBERO and what happens if the same registry do repeat same emergency threshold?

CHRISTINE WILLETT: Just in general, I know that there are a number of questions about the specific procedures around EBERO so it might be helpful if I could point you to the common transition process. It's a portion, the EBERO agreements with our EBERO providers are published on our website and at the end, Exhibit B of the version I'm looking at here, it includes the common transition process. So it has a lot of detail about these processes and procedures about who's involved in the decision making from an ICANN perspective, this multi-functional, cross-functional team that the organization puts together when invoking an EBERO process. So that might be something you want to look at after today's session.

I don't know. Francisco, do you have a more specific answer to the question?

FRANCISCO ARIAS: No, I don't think so. I will only add this multi-area group that looks at the issues when they are happening in the technical side, the security side, and – sorry, what's [inaudible]? SSR. The SSR [area] is [inaudible] we get to those discussions, we have executives like Christine, and so it's multiple people that are involved in such a decision.

JAMES GANNON: Just a quick question. Is there a nominated decision maker within that group who holds the final say on go/no-go?

FRANCISCO ARIAS: So the final say, like anything else falls in an executive, particularly in the case of the EBERO, it will usually be Akram as the head of GDD. But this is [inaudible] by 24 things that could happen anytime. We [can] rely on other executives to make such decisions.

CHRISTINE WILLETT: Correct me if I'm wrong here, Francisco. If I could just add a little detail, this is an executive steering committee that guides and is consulted with, but there is an event director. We call them the EBERO Event Director and in the process documentation, you'll see this role called out. We have, I want to say, three trained, anytime – four now? – trained event directors now. So that is the person who leads this ad hoc team that comes together for the individual event and then they consult with executives, including such as Akram that's always an officer on that group.

BOBAN KRSIC: Just one question. Is the current registry operator also involved in this decision or is he out of scope?

FRANCISCO ARIAS: No, they are not part of the decision making.

NOORUL AMEEN: Yeah, just to clarify one more thing. I ask about cases when [inaudible] is repeating for the same registry. You mentioned before that it is cheaper and better for end users just to let registry to fix the problem than to move to EBERO. But if that happens so often, what's going on?

CHRISTINE WILLETT: So there are cumulative thresholds for invoking EBERO and invoking the transition, not just 10%, 25%. Francisco, can you address, it's a monthly cumulative threshold?

FRANCISCO ARIAS: Right. So in terms of the compliance side, service level requirements are on a monthly basis. It's [inaudible] during the calendar month.

The emergency threshold, it's measured differently. It's measured on a seven-day window, moving window. The last seven days are taken into account to consider, for example, the amount of downtime that has happened during that period of time.

JAMES GANNON: Have you ever had an RSP that has reached a cumulative threshold?

FRANCISCO ARIAS: Yes. That's what the [inaudible] are. They are cumulative thresholds and the compliance thresholds, but they are even lower so there have been more [inaudible].

NORM RITCHIE: Another random question, I was just looking in the notes here and I noticed we didn't talk about the escrow at all, and the EBERO process actually relies on data escrow, correct? [inaudible] the data comes from. So how do you know the data escrow is fulsome? How do you know it's accurate? Are there tests that are done or something?

FRANCISCO ARIAS: So I would like to clarify something there. We would rely on that escrow on the worst case scenario in which you have raised that is cooperative because that's not necessarily the case. They may be having an issue and they still want to be good citizens and help in fixing the problem. The worst case scenario in which they are either willing or unable to provide help, yes, we will go to using the data that is in the data escrow deposit. And not only that, it's also the [inaudible], which ICANN receives directly so we have a process to check the [inaudible] daily to ensure we are getting some data, and so we will use the [inaudible].

In the case of data escrow, what we have is the [inaudible] requirements to provide deposits on a daily basis and they have requirements to report to ICANN, so we receive some high level summary data of the data that is being received. There is also a verification process that has to be executed by the escrow agent and they also have to report to ICANN. So we get that data and we will find

that they are compliant with their obligations, and that's the kind of checks we are doing on [inaudible].

NORM RITCHIE: No, but what I'm getting at is how do you know that data is actually accurate? Or let's say I send in a deposit to the escrow and I send half the domains, and they said everything's fine?

FRANCISCO ARIAS: We don't have access to the data that is in escrow, so we cannot do any checks on that. We are relying on that verification process that the [inaudible] are required to do in order to have the best technician possible on the data.

When that data is received by the EBERO provider as part of the process, there is a checking process. They find the data versus the zone file. That is something, for example, we have to [inaudible] the difference and they [help us] here. I think it is in the [CDP] document that Christine mentioned. I don't know if it's in that or internal procedure. But it's the checking of those two that are sources and there are [inaudible] that are directions there on what to do about those cases and that's what we have.

DENISE MICHEL: So I hadn't looked at the ICANN coordinated disclosure guidelines in quite a while and these are the guidelines that were issued in, I think, 2013 to lay out what ICANN's role is when vulnerabilities are identified and ICANN determines that these vulnerabilities are an SSR risk and

could be exploited [inaudible] and they also explain steps that can be taken.

I believe the SSR Team had the lead in implementing this coordinated vulnerability disclosure reporting and sort of coordinated action. Does your team have any involvement in this, or has your team used this disclosure reporting system to report any of the issues that have arisen that we've talked about today, any intersection there at all? Thanks.

FRANCISCO ARIAS: So in [here] it will be users of that procedure or as part of the ICANN organization, if you like. We don't have a role in [inaudible] SSR [inaudible] and I don't know who else in defining that process. It was, of course, presented to us like any other [inaudible] within ICANN.

I don't see much intersection with either process or the [SLAM]. We have used that process in other cases that are public like the [inaudible] was in contrast, part of the [name collision] issues. So for that, we followed the process that is not related to the SLAM or EBERO.

DENISE MICHEL: So there hasn't been [inaudible] an instance where a vulnerability that you feel has posed a broader threat that has come up through your work or related work and you haven't provided any input into that system. You see yourself as more of a responder and have not actually added anything into it in terms of a threat or a vulnerability. Is that right?

FRANCISCO ARIAS: I guess so. As I mentioned before earlier in the session, we have found cases where the [inaudible] were caused by [invoking] their system, so they are the ones that could report and we were told they reported to whoever developed the software to fix the issue they encountered.

We're not the end operators of the TLD. We didn't have the details to go and talk with the [inaudible], so no, we haven't had a situation in which we needed to use the [inaudible] process.

JAMES GANNON: This is probably, actually for Steve. Without details, do we know how many instances in which the CBD policy has been inactive since it was put in place?

STEVE CONTE: I could say certainly one, but I don't have further details around that. I was just actually scouring our website to see if we actually did any kind of public reporting close to vulnerability and quickly, I'm not seeing anything but I can go back to my team and ask the question, and see if it's acceptable to answer. But I know we've talked to him publicly before that we've had at least one vulnerability that was held for a while between organizations while that vulnerability was mitigated.

UNIDENTIFIED MALE: Okay, it's a random question, then I would say thank you all for your time.

STEVE CONTE: I do have the process for the vulnerability link if you guys want to post it into the chat if it's helpful at all. I'll put that in.

UNIDENTIFIED MALE: Thanks, Steve.

FRANCISCO ARIAS: So then [inaudible]. Our next session is an internal session or [inaudible]?

JAMES GANNON: In particular, thanks from me for standing up to all my grilling questions on EBERO.

FRANCISCO ARIAS: Not only yours.

DENISE MICHEL: Yeah, thank you very much for sharing your time with us. It's been really helpful and we've also submitted questions in writing. I think that staff is working to get answers back in writing and as we continue our work, I'm sure we'll be sending over some additional questions. Thank you so much for all your help with this.

NEGAR FARZINNIA: Correct. We covered the domain abuse earlier. So Boban, back to you.

ERIC OSTERWEIL: So I don't know if they can page someone else, but we could go over our planned agenda list of things that we wanted to talk about and see how many of the items that we're curious about we touched on, covered, closed out, whatever today, right? Check boxes. Does that make sense?

UNIDENTIFIED MALE: I got [inaudible] issue. Can I contribute one more thing to DAAR before you guys go into there?

During the break, I was able to confirm because the public agenda is not posted yet. There will be a public session on the [inaudible] abuse reporting in Abu Dhabi on the 31st which I believe is Tuesday, I think. Absolutely we'll send out a list of suggestions for the Review Team as a plenary, but I just wanted to confirm that there will be an ICANN discussion about the abuse reporting tool, how we use it and how it can be used by the community.

DENISE MICHEL: Can I ask a question? So could I ask this be recorded as a question to have answered by whoever is most appropriate in writing back to the team as soon as it's appropriate? I think the team would benefit from hearing from staff, like an e-mail from staff, that discusses the different challenges that you noted that were inherent in conducting, gathering abuse-related data. So I guess this is DAAR specifically. Is that what you call it? DAAR?

Yeah, and so, could staff give us an account of the various challenges of gathering and publishing the DAAR-related data?

STEVE CONTE: Absolutely. I will pass that on to my team and we'll respond the best that we're able to and within the constricts of some of the contracts that I had spoken about earlier that we might not have full ability to speak about some things.

DENISE MICHEL: Sure, and do feel free, of course, to note if there is some confidentiality involved in providing a full [term] response to this question.

STEVE CONTE: Sure.

BOBAN KRSIC: Okay, thank you. So back to today's agenda and to the item, business continuity management, that we had this morning where we talked, I would say, 70% of the time about risk management.

UNIDENTIFIED MALE: 90.

BOBAN KRSIC: 90? Okay, I would say 70. And let's identify the items that we have or that we can maybe address tomorrow because we have Xavier and

James tomorrow at 3:30 for one hour for risk management slot. And I think there are some open topics around business continuity because there is only, we have some procedures in place but maybe it would be good to see something documented and to have a look inside and to go deeper in these topics.

Yeah. No, that's [raised here] but we will identify some issues or some items from the business continuity and move them to tomorrow's agenda.

UNIDENTIFIED MALE: [Inaudible]

BOBAN KRSIC: Yeah, it is. Yeah, but tomorrow we will have one hour as well.

JAMES GANNON: Yeah. So if Negar or Yvette could give a heads up to James and Xavier that we'd like to do a little bit more of a deep dive into business continuity tomorrow rather than the risk treatment which I thought we covered today because I feel like we got the topics switched around. So if we can give them a heads up that particularly any procedural documentation that they're able to surface to us without NDA, that they're happy to make public now and that will be extremely useful and even any high level statistics or some data to back up some of the discussions that we had this morning, that that's what the team would like to look at tomorrow.

NORM RITCHIE: Sorry. Yeah, specifically, it says how it relates to DNS because that's a very, very broad topic, especially for ICANN. It's a big organization but this is about DNS.

ZARKO KECIC: Yeah because the last item tomorrow is AOB so we can expand that if you give us opportunity and room. Can you ask them to think about extending that one hour to two hours or an hour and a half?

NEGAR FARZINNIA: Thanks, Zarko. I will check with them. I'm not sure what their availability is like, but I will definitely ask.

ZARKO KECIC: That would be nice because this is short and today was short as well.

NEGAR FARZINNIA: Well, yeah. I'll definitely propose that to them. Thank you.

ZARKO KECIC: Or we can start earlier because we have a subteam face-to-face discussion time slot beginning at 3.

STEVE CONTE: This is a logistics question. Does anyone have an early flight tomorrow? What time is your flight?

JAMES GANNON: I will have to leave here at 3:00 tomorrow, unfortunately.

STEVE CONTE: Okay. Anyone else have a flight? So you're asking to extend it, that will push us over, assuming that we're staying on time tomorrow, it'll push us over time because they're the last group to meet with us. I just want to make sure that we had enough people here that were still going to be here if they agree to go over. So other than James, it sounds like everyone is still here for enough time to have a conversation, correct?

ERIC OSTERWEIL: I may turn into a pumpkin at 5:30ish.

STEVE CONTE: Happy Halloween.

BOBAN KRSIC: Okay, anything else from today's agenda to address tomorrow? We talk a lot through the operations stuff and systems related to them. I think we have a good overview and I'm fine with this.

Okay, that's it. Then only business continuity stuff and extend tomorrow's time slot with Xavier and James.

ZARKO KECIC: And to try to get more precise answers from them.

BOBAN KRSIC: We can try.

ZARKO KECIC: [Just to] push.

BOBAN KRSIC: Anything else?

JAMES GANNON: No, just that I think we covered a lot with GDD and I think that was useful, particularly I think we fleshed out a lot of the concerns we had around EBERO. I can certainly see one or two recommendations coming out of those discussions. I have some specific things that I think I'll follow-up on and potentially start drafting some language around and I'll probably write it while I'm on the plane tomorrow. So I thought that was a very useful session, and again, I'd like to go back to Francisco and I've forgotten the lady's name already, Christine, that we felt that was a very useful session. I think it was good to see the openness around the more tricky questions. I think that was good.

DENISE MICHEL: I was thinking it would be helpful for the subgroup, before we disband, I didn't mean to say that, I'm sorry. I'm getting a little tired. So yeah, before we depart L.A., just to sketch out what we saw as the key points that we gathered more information on and a cut and paste of our agenda with some additional detail and if in the course of these two days, we come to an agreement on further refining some of our issues, adding or dropping any issues, I feel given some of the external comments we've gotten about this fact-finding two days that we're spending here with staff, that it would be good to have some semblance of a report on what we did and where the subgroup work stands and if we're able to get an agreement on that before we leave L.A., I think it might make it easier for us. Thanks.

BOBAN KRSIC: Maybe we can start tomorrow at 10. We have one and a half hours for day-to-day discussions to wrap up this day, yeah, and to just write it down, the main points to keep them in the mind. So yeah, good, perfect. Anything else?

If not, I would say thank you all for your time. Thank you, Steve. Really good job, yeah, and [inaudible].

I start with Steve on my right side, and also Yvette. So yeah, see you tomorrow, yeah, at 9:00. No, 8:30, a meeting and we have –

[END OF TRANSCRIPTION]