STEVE CONTE: Okay. Good morning, everybody. It's Day 2. I tried yesterday, but they wouldn't move the restrooms, so the restrooms are in the same spot as they were yesterday, so I really don't have much housekeeping to do today.

Negar, do you have any announcements?

NEGAR FARZINNIA: Just one thing. I did want to let everybody know that Ash actually wanted to stop by and meet the Review Team and talk to everybody, but unfortunately, his house in Orange County is in evacuation path because of the fires, so he has been preoccupied with that. He is not in the office, unfortunately, and we don't know when he's going to come back. So he sends his apologies and hopefully in the future, he'll get a chance to touch base with everybody and at some point meet with the Review Team. I just wanted to pass that along.

STEVE CONTE: We did, at least up until a minute ago, have Darren Kara with us who will be back, but he's not scheduled here until 9:15, so Boban, that's all the housekeeping I have. I'll pass it on to you at this point.

BOBAN KRSIC: Hi, good morning, also, from my side. No, nothing. Only for the review of today's topic. So we have, now at 9:15, ICANN Security and topics that are related to DNS Security. Then at 11:30, we have a slot with

WHOIS Compliance followed by security incident management, and after a break, we will have topics from yesterday, [inaudible] related topics, and risk management from today's agenda with Xavier and James.

No, that's from my side and yeah, we are waiting for Darren to start with the first topic. James?

JAMES GANNON: Good morning. I need to give my standard legal disclaimer once again that I'm here as an individual, not representing my company. Any opinions or anything that I say here are my opinions and not intended to represent my employer. Thank you.

STEVE CONTE: So with us this morning is Darren Kara. See? I almost mispronounced it even after I talked to you. Darren is relatively new to ICANN, but I think an old soul in the community and he's our DNS Engineering Strategy Manager and will be talking, having discussions with you guys about the SSR of infrastructure and things like that. So Boban, I pass it to you and to Darren and however you guys want to run it.

BOBAN KRSIC: Hi, Darren. Good morning. Yeah, as [he mentioned], we would like to talk with you the next half an hour about ICANN security procedures, processes you have in place that are related to DNS, and [inaudible] of DNS. So that was the main point.

We organize it that we are going to a dialogue that you start with and a broad approach to talk about the processes that you have in place, and then we will [go on] dialog and pick up some main points out of them and discuss them.

DARREN KARA: So the general question is the ICANN root server operations, otherwise known as L-root also amongst root server operators. We follow all the standard DNS procedures set forth by the community in RFC 7720, as well as the RSA CC01 document. We've answered those and published those as far as the processes [befall] for a robust DNS infrastructure with which we support the root server operations and the root zone.

We have a diverse team, geographically dispersed around the country and the world, so that we are not in the situation where if we lose a site, we lose the management and resiliency of the root servers.

We have about 170 root servers hosted around the world, as well as three core sites that are clustered with which we handle root server traffic for the L-root.

Host companies are vetted through a process of a checklist for technical security as well as physical security so that we know that L-root systems are in good locations and we have processes for placing those servers, as well as removing those if we feel that maybe local hosts have no access to L-root. It's only managed the DNS Engineering Team within ICANN. No other ICANN employees have access.

We manage all of the instances from two sites in the United States, one on the East Coast and one on the West Coast so that if we lose either the East or the West, we will still be able to manage the root zone infrastructure, make sure the root zone gets signed and published and is readily available for the Internet At-Large.

We work closely with the other root server operators in order to mitigate larger attacks upon the community. There are tabletop exercises that occur every year to discuss things and to role play what happens during an attack or a compromise or [inaudible] exploits for DNS applications.

We also work closely with the DNS vendors so that we are ahead of the curve for security response. They generally let the root server operators know if there is a possible issue so that patching can be done before a general announcement to the public, which helps the resiliency of root server operations overall.

We do constant monitoring. Not just us, internal to ICANN, we monitor all the root server operations, the signing operations, and the zone transferring. But we know that things like DNS [inaudible] from RIPE are constantly keeping statistics on our availability as well as Verisign. Verisign very robustly monitors all the other root server operators, so we have multiple ways of getting feedback as far as our performance and how we look on the Internet as a whole, serving up DNS.

I think that's generally my opening statement for this. If you have questions specific to any of those areas or if there's something I missed, let's get into it.

STEVE CONTE: Perfect, thank you. Denise?

DENISE MICHEL: Hi. Thank you for spending time with us this morning. So one of our jobs is to review the implementation of the first Security Review Team recommendations that the Board approved and directed staff to implement almost six years ago, assess not only their implementation, but also their impact.

Recommendation 9 of that report says that ICANN should assess certification options with commonly accepted international standards for its operational responsibilities and publish a clear roadmap towards certification.

So the staff report in stating that it fully implemented this recommendation refers to SysTrusts, now SOC audits, of individual IANA functions. Of course, those SysTrust audits were already in place before the SSR1 read its report, so the team clearly wasn't referring to that.

The staff report also indicates that ICANN has incorporated SSR-related certification in its EFQM program. We haven't received any information, I don't think, about these audits and we were also told that some of the engineering staff was ITL-certified. That seems to be done on a case-by-case basis, so it's not obvious how ICANN assessed the certification options as a result of the first security review report.

I'm fully open to the possibility that I've missed something, so please. I'm just going to lay this out and I really invite you to correct me and help fill in what's likely knowledge gap.

So it's not obvious how ICANN is assessing certification options or did as a result of SSR1. Why ICANN didn't publish a clear roadmap towards whatever certification it thought was appropriate – again, perhaps I missed that – and what the current state of certification in this area is. I'll stop there.

DARREN KARA: I just want to see, I just want to clarify a couple things. First of all, the SSR1 was the same as SSR2 in the fact that it was a security and resiliency review of the unique identifiers in which ICANN manages and has coordination efforts on. It was not a security review, necessarily, of ICANN, ICANN infrastructure, or anything like that. It's as much security as this Review Team is looking at.

So I know that was probably shortened for brevity and was talked about yesterday when you brought up the security review, but I just want to clarify for the record that it was a security, stability and resiliency review, exactly like what you guys are having right now in a Review Team.

So it sounds pedantic, but there's a fundamental difference between a security review of ICANN infrastructure and ICANN audits and things like that versus a community SSR review.

DENISE MICHEL: Yeah, I'm not quite understanding. Does that have something to do with the questions I've asked?

DARREN KARA: It does, in some sense, I believe because a security review would be something that, I believe that we would contract someone to do, something like a pen-test or other review to make sure that we're robust in our security efforts. This review is talking about the security, stability and resiliency of unique identifiers in which ICANN manages and ICANN coordinates. I think there is a fundamental difference between that, especially in the information that can be provided around the two different areas.

DENISE MICHEL: Thank you. So getting back to Recommendation 9 of the Review Team that the Board accepted and directed staff to implement, can you elaborate on the certification options that were then pursued and whether a roadmap was laid out for what ICANN staff assumed was the appropriate certification and standards to be followed for its operational responsibilities?

STEVE CONTE: The various types of certifications and things that are mentioned in that span beyond IT. You mentioned EFQM and things like that. I think this is a question that's best to be captured and responded back with an e-mail and a written response on that so we can make sure we get the

right answers to the people who are managing those various levels of certification. James?

JAMES GANNON: Okay, so I'll simplify the question with him. So within the scope of your responsibility, how do you manage deployment of new instances, change management? Do you do an ITIL style process? Do you have an internal process that you do to do hatching, to do security management, vulnerability management? How is your change management eventually done within the scope of the DNS engineering and the L-root?

DARREN KARA: We do constant patching. We have a robust puppet infrastructure and depositories for all of our systems. We do a puppet run every 30 minutes to make sure that the proper configs are in the proper places. Twice a week during maintenance windows, we do our typical updates for general patching of systems as well as upgrading software. We test that in staging environments so that we do not cause any hits to our root server operations.

If there is out-of-cycle changes, we have a group review for the change. We make sure it gets tested, we make sure it gets vetted by each of the team members, and then we can push it off-cycle if we think there's a security concern or stability concern. We do this pretty regularly when it comes to DNS software.

Along those lines, we have the ability to switch between two different name server vendors on each of our hosts if we think that one of the

name server domains either has an issue that's going to cause issues with stability or security, so that we can flip to the other vendor. We split it 50/50 on our hosts right now.

All our hosts, it's all key-based authorization for connecting to them. It's a push from centrally-located systems in East and West. All the hosts have individual firewalls on them, which keep them [ACL] from any traffic that is not from our control structures, other than, obviously, the traffic that's needed for DNS and zone transfers amongst the Internet.

UNIDENTIFIED MALE: Just regarding to this, all these procedures and controls that you mentioned, is that documented somewhere?

DARREN KARA: Yeah, we have our internal documentation and the [get repose] also, replicated East and West so if we have issues with one site, we have access as employees to our documentation in order to follow our best procedures.

UNIDENTIFIED MALE: Boban, thank you. You're running the queue.

BOBAN KRSIC: I'm running the queue. Okay, then, James or Eric.

ERIC OSTERWEIL: So do you guys do any IDS or IPS on your root server instances to detect anomalous behaviors, intrusions, anything like that? Do you have [inaudible] posture on the end nodes besides firewall?

DARREN KARA: We ship logs off for central log correlation, but due to the nature of where the hosts are, we don't have access to a lot of the network infrastructure in order to watch the flows coming in and out of each of the root servers.

ERIC OSTERWEIL: Yeah, thanks for that. So then, follow-up question: are the hosts in container management? Are they virtualized or do you run them on [verometal]?

DARREN KARA: No, these are all [verometal]. Once somebody comes and requests to be a host, we point them at certain specified hardware that they can purchase that if they meet the criteria and they landed in their data center, wherever they have it, we then take control of it after it has basic track functionality and then we [inaudible] onto it.

ERIC OSTERWEIL: Okay, so you're the only application running on that system.

DARREN KARA: Yes. It is dedicated for our use only.

ERIC OSTERWEIL: Thank you.

DARREN KARA: No shared environments, no virtual environments.

JAMES GANNON: Thanks. Two quick follow-up questions. So it sounds like there is a lot of custom process, but a lot of internally-derived processes around how you're managing a lot of this. Is the reason why you haven't looked at a more industry-standard process, which is ITIL or something less unique in line with the SSR1 recommendations to move to a more industry-standard, certifiable and auditable method around [inaudible] founding controls, vulnerability management?

DARREN KARA: So coming into this, in the process where I have, I'm not sure on that front. I know it's a homegrown solution that follows, from my perspective, pretty general guidelines as far as repositories and pushing changes to multiple systems. I'd have to get back to you on the ITI help and if we're going there or why we haven't gone there.

DENISE MICHEL: Just to make sure we capture this, could staff capture the questions that both now I and James have asked, so when you have a chance to respond in writing, we can just make sure we document that?

STEVE CONTE: We are recording this, so we'll be capturing. There was a lot of dialogue yesterday and there will be a lot of dialogue today and we are capturing. We'll be going back over the recordings and making sure that we capture the questions. As I mentioned yesterday, we'll send them back to the Review Team and make sure that the accuracy of the questions before we go back and answer them.

BOBAN KRSIC: Do you want to add to that? Then we have Norm.

NORM RITCHIE: I believe you mentioned that you're running two flavors of DNS software. Okay, now what about the OS?

DARREN KARA: We have two different flavors, primarily one, but we are doing the development work in order to have an equal spread of two different OS flavors and avoid the same [inaudible].

NORM RITCHIE: Redundancy through. The other thing I just wanted to note, actually, is right now you said there's like 174 nodes, I believe. What was it five years ago when SSR1 was done? So there's been a big change. There's been a big growth in the area, right?

DARREN KARA: Yeah, I'd have to get back to you on that. I wasn't here for it. I don't know the growth curve as far as pushing out our single instances.

NORM RITCHIE: Yeah, I was actually going to commend you on it. I think it's a great thing that you've expanded it.

DARREN KARA: Oh, you know the original numbers.

NORM RITCHIE: Well, no. I know it's smaller. I don't know what it was.

DARREN KARA: I believe as far as root server operators go, we have the largest install base right now as far as the named letter groups.

NOORUL AMEEN: Okay, well, I just noted in here. There's a particular RFC for the root name server operation requirement that is RFC 2870. It tells all about the security requirements for a root name server implementation.

STEVE CONTE: For the record, we're just pausing that question because we had a spill here.

NORM RITCHIE: Is your responsibilities, are we talking only with L-roots here or is it [zone] DNS as well?

UNIDENTIFIED MALE: Other than ICANN?

NORM RITCHIE: Yeah, so I'm wondering about IANA. How is the DNS managed for IANA.org? Is that in your baileywick?

DARREN KARA: I'm just trying to understand the question better and I'll summarize some of what may be the answer or may not be the answer and hopefully that'll start that up.

So the IANA domain still resolves, I believe, on some of the ICANN-operated servers, but I'm not sure if that is the question you're asking.

NORM RITCHIE: Well, given the relative importance of IANA.org, because if you have to reach out to it to do any automated changes to a TLD zone or whatever, you have to be able to reach it, so I'm just wondering how the DNS, what the resiliency is for IANA.org to maintain? Is it given any special attention, I guess, is kind of the question.

STEVE CONTE: So I'll give you a chance to answer this as well, just doing a quick WHOIS on IANA.org and NSSEC@ICANN.org is one of the four name servers. The other name servers fall into the domain of IANAservers.net, and unless Darren has an answer or information about that domain, that's something we would have to capture and take it back to PTI. Darren, do you have any further info?

DARREN KARA: No, other than if it spread across four different name servers and one is us, it certainly falls under the same robustness and resiliency as any of the zones we serve up.

NORM RITCHIE: I will note, too, that you have registry lock on it, which is great. I'm glad to see that.

DARREN KARA: James [inaudible], I will capture that, though, and we will provide an answer on who administrates IANA-servers.net as well, and any process behind that.

STEVE CONTE: So do we want to go back to the previous question about RFC?

DARREN KARA: I want to wait until Noorul is back because he was the [asker of the] question.

JAMES GANNON: I have a follow-up in the same general areas we were just talking about, so also expanding just away from just the IANA.org stuff, there is a lot of what I would term "critical operations stuff", the GD portal and some of the other supporting infrastructure. Again, I think we can probably add that to the written response question of, is that managed just as a general ICANN service or is actual attention given to the supporting infrastructure that's [inaudible], that sits around the root zone maintenance and the [inaudible] season? And how is that handled or is that just handled as a general ICANN IT service?

DARREN KARA: Back to you on that. We do have a list of zones that we pay particular attention to. It's easily pulled from a file, so we can respond to that in an e-mail.

ERIC OSTERWEIL: Yeah, so following some of these questions, train of thought, so have you guys mapped out sort of the systemic dependencies that are involved to maintain the L-root deployment, whether it's you get [inaudible] instances and what services and systems are needed to have that operate so you can do your puppet push and stuff like that? In other words, what are the systems that keep L-root alive and up to date? And where is that maintained and can we see or is there some way to get an idea of what systems are involved with keeping that operational? Not a serving port 53, but keeping the instances in the [inaudible].

DARREN KARA: Maintaining the business that we maintain. Yeah, as I said earlier, we have the infrastructure in two sites so that if one were to go bye-bye, we would still have access from the other site. And that's a staging environment, get reposed management servers that are ATL and have access to the different instances so that we can continue to push updates for OS and name server software.

We have key signing infrastructure locked in safes in both of the co-location facilities so that we can continue securely signing zones through the zone signing, and we have flow charts for this in internal documentation that show each step as far as checking in changes, reviewing the changes, doing things like [inaudible] for some of the testing before it gets pushed out.

ERIC OSTERWEIL: That's awesome. That's great. No caveats. That's really awesome that you guys track that. I just wonder one step further, do you track, sort of the secondary dependency is like when there is a domain name that needs to be looked up because GitHub.com or your internal whatever naming infrastructure, does it reach outside whereby at some point, if there was something going on with the root and one of your dependent systems, have you done even just a [packet] trace level or any kind of audit to see what sort of expected or unexpected external dependencies might keep in the case of an event, disaster, or whatever else, keep you from being able to maintain your deployment? Do you know what I'm saying? When you do the tabletops?

JAMES GANNON: I can give, probably a specific example. For example, I know ICANN uses OPTA a lot, across a lot of your systems. Do you have where you're logging in to your GitHub instances through OPTA and you have backup procedures to be able to fall back to local user, these type of things where you're not mapping necessarily the direct infrastructure, but the supporting infrastructure that is around it and thus, is impacting your ability to use that supporting infrastructure.

DARREN KARA: In operators. We do kiosks for all of our systems so that as new members come on board, we supply a public key and it's distributed via puppet to all of our systems so that we can directly connect and do our pushes and changes.

We are not dependent on ICANN's larger infrastructure for most of what we do. Ideally, if we're dealing with strategy and management, which I was brought in to do, I would like to make sure of that. If I got that mandate to make sure that I could pick up one of our sites and move it somewhere else, set it up and be able to have access too. I think that's the case right now, but it's a good exercise. I'll definitely take that back.

STEVE CONTE: Noorul, you were questioning about RFC 2870, I believe.

NOORUL AMEEN: So it talks all about remote logging should be disabled and physical [degradation] be enabled, those kind of [inaudible] are there. So my first question is that how you're verifying the organizations which are implementing these [inaudible] name server operations? They're ensuring this RFC of compliance.

The second part is that we have served many illegal root name of server operations in many of the countries. What is ICANN's stand on those kind of implementations?

DARREN KARA: So just case in point with RFC 2870, I believe is obsoleted by 7720.

NOORUL AMEEN: Yes.

DARREN KARA: Okay, and we follow RFC 7720. We do not allow the operators to log in and tamper with our systems, so we turn off root login. They cannot get consult.

NOORUL AMEEN: What about IT because many of the systems, I see [inaudible] enabled and in the world of protocol abuse, in future I think may be abused to launch any kind of cyber attack, so what's the stand on ICMP enabled? It's good for tracing and for echo purpose and all the things, but still.

DARREN KARA: To each of the servers – what? – allowing ICMP packets in?

NOORUL AMEEN: Yes.

DARREN KARA: We have it enabled, I believe. I have to check.

NOORUL AMEEN: If you think these are all server IP addresses, you can steal lots of packets at a time.

STEVE CONTE: Speaking from ancient history because I used to run L-root a long time ago, but we used to keep – and this is eight or nine years ago, so I can't speak to today – but just from our head space from back then, we used to keep ICMP on because, frankly, the attacks that we would see on the root servers back then were using DNS and DNS was a larger packet and a larger payload. So the ICMP, even if they're trying to DDoS us, it was relatively on a smaller scale than some of the reflector attacks and things that we would get from a DNS-related attack.

So that's what we used to do. I know ICMP is still on because a lot of probing, like the RIPE analyst and stuff uses not only DNS queries, but ICMP and Ping tests to see the live status of a server. So I presume that based on my experience back then that the same mentality is, they're now with the exception that they even have a more robust

infrastructure and larger pipes and bandwidth. So as far as from an attack perspective, it's probably less of an issue even now than it was back then.

NOORUL AMEEN: See, five years back, we also thought like that, that protocols like NTP, SSTP and [this ICMP] won't get abused for launching large [inaudible] reflected kind of [attacks].

But all together with this ICMP and any other parameters or [inaudible] or any kind of amplification, all together, will it invite any sort of attacks or something like that?

DARREN KARA: Since joining, we haven't had any attacks since at least a year before I joined the team. If we feel that for resiliency sake, that one of our hosts is being abused, we will withdraw the announcements so that the larger L-root infrastructure doesn't have problems in general. So no, I don't believe it's currently a factor of attack. DNS is definitely the deflection that people attack.

NOORUL AMEEN: Second part, illegal root name server [inaudible].

[STEVE CONTE]: Can I take that? That's a term that we don't ever use in ICANN. We don't' necessarily view any. Anyone can run a DNS instance. Anyone can

run their own root. ICANN and the ICANN community supports the IANA authoritative, what we in the community are calling the authoritative root.

But there's been alternate roots since the dawn of DNS. There's no mandate within ICANN or with any team within ICANN to say we must snuff those out. There's value to be had about running alternate roots. We, as a DNS community, learn from the different ways that alternate roots are operated.

We've seen everything from ISPs. Earthlink, I know, way back when used to run an alternate root. The various members of the DNS community run it for profit and for research, whatever they want to do on that. There's nothing that we consider to be an illegal root, even if someone is populating the same data that's within the authorized IANA root, even if it's not illegal, as long as it's not creating stability or resiliency implications of the IANA authorized root, then people can do what they want on the Internet.

One of the things that will help stem or help solidify the authoritative root of IANA, though, is the DNSSEC in the DNSSEC [tree]. It's hard to replicate any instance of the IANA root without having the trusted chain on that. So Eric, you're probably going to correct me on some of my terminology and I welcome it.

ERIC OSTERWEIL: No, no. I'm actually going to just try and reframe the question more generally to sort of marry the two perspectives, I think.

So without getting into who's an authorized root and who's not an authorized root, who's a good root, who's a bad root, are there any checks for consistency run across any of the deployments of the data of these instances starting with the ones that you manage, but nominally if you wanted to broaden it, you could say, do you check the different points of service of a reported root are consistent.

I imagine you may not want to make it your business to go into other people's ballpark, but in your own, [inaudible] on deployment, do you have any sort of measurement apparatus or statistics on how consistent and what the replication?

DARREN KARA: We measure the time for our zone transfers when changes are made and also do comparison against what we know to be the master root server, so if there are any changes, it raises a flag as far as that host goes and it's something we can take action to remove the host.

And actually, that's done automatically. If there's too much adrift or too much change between one of our systems and what we know to be the proper root zone, it'll withdraw the announcement.

ERIC OSTERWEIL: So that's really cool. That's awesome. Are those statistics available to the team, or the community or the public? Are they internal, especially longitudinally? The root has changed a lot and certainly those statistics become useful measurement points for other explorations, but then, and just to sort of add two questions into one breath, unfortunately,

sorry, I'll forget it if I don't say. Do you also check the data consistency in addition to just AX or IX or whatever you're using? In addition to transferring the zone, do you actually check that it's serving the right values?

DARREN KARA: Yes, we do. As far as supplying the data, we can see. If you come up with specific data sets that you're interested in, we could see if we can capture that.

ERIC OSTERWEIL: It was a general question, but yeah, I might actually follow-up with you. Thanks.

STEVE CONTE: Just to clarify that, do we need to capture that within the Review Team or is that going to be more of an interest level between the two organizations?

ERIC OSTERWEIL: No, it would be on the Review Team. I think what we would do is we, as a Review Team, decide if there's something interesting for us to look at from an SSR perspective and if we could use the data for something beneficial to the Review Team, we would ask for it.

STEVE CONTE: Okay, thank you. I just wanted clarify how that interaction would take place. Thanks.

JAMES GANNON: Sorry, I'm going to pull you back from policy and process. Security incident management, do you follow ICANN's general incident management process or do you have anything specific to L-root on if there is a vulnerability identified out in the market that may have an impact on L-root that you have your own incident management process on how you handle that with regard to, it might be emergency deployment of patches, etc., and is that a standard documented process that you have within the L-root team?

DARREN KARA: We can do changes outside of our typical change management and we keep a specific eye out for anything that could affect root server security. Yeah, so it is under a process because we know we need to be able to act quickly. And sometimes the process is, as I stated earlier, we know there's an issue with name server X, okay, we are going to switch over to name server Y until we know that the community has it fixed. Or, a lot of times, like I said, we are identified early about zero days and things like that, and we get the changes before other people, we can then retroactively announce to people that we have addressed a vulnerability concern and taken care of it. We don't want to send up a flag, saying, "Hey, we're vulnerable; come after us."

STEVE CONTE: So this is possibly more. Just to state for the record that the session immediately after lunch is also with Gary Petzer and Sam Suh, and I believe Darren might be back, but that whole session will be about instant response, or security incident management.

JAMES GANNON: It may fall over into that.

STEVE CONTE: Yeah, but I just wanted, before we get too deep into that.

DARREN KARA: Yeah, I'll make sure that it's [inaudible] too.

JAMES GANNON: [Inaudible] response but it may fall over into the next session.

So with regards too, again, your domain, DNS and the L-root, how do you monitor for potential vulnerabilities? Is it just communities and person-to-person knowledge or do you subscribe to threat intelligence feeds? And feel free to be generic where you have to be.

DARREN KARA: We use a third party vendor for scanning as well as have internal tools that we use to scan so that we have two vectors to analyze how we look from the outside and from the inside. So yeah, we are constantly scanning our own systems in order to not just watch out for DNS

vulnerabilities, but we get reports of system-level, down the file system of where things that we might have issues with permissions and things like that. So we have reporting going on daily for that.

DENISE MICHEL: I don't know if this is appropriate for this or the next set of staff, but could you speak a bit about the status of things? Several months ago, some DNS attacks and defense occupied a lot of attention with attacks taking out some of the major root zone servers. How do you feel the state of things are and can you elaborate, if not now, perhaps in writing about steps that were taken, both by ICANN and by the community in response to that?

DARREN KARA: DNS or the general state of root server operations?

DENISE MICHEL: I'm referring specifically to the attacks that jeopardized some of the root servers and the community response to that and ICANN as well.

DARREN KARA: I think we would have to do a written response to that question.

DENISE MICHEL: Thanks, I'd appreciate it.

STEVE CONTE: As the root server operators have ascertained for years and years and years, each root server operator speaks only for themselves in the sense of that too, so I think a written response would be ICANN's perspective as a root server operator to the sense of things and not speaking as RSSAC or the root server operators as a whole.

BOBAN KRSIC: We have two minutes left. Do we have questions that we can address now, or one? Okay, the last one.

NOORUL AMEEN: Is there any meta data analysis happening on the DNS data or any global DNS log monitoring happening to prevent DNS level of attack or something like that?

DARREN KARA: Restate the question.

NOORUL AMEEN: The first part is there a meta data surveillance part in the DNS traffic or something like that? And second part is the DNS log monitoring. A DNS log will tell you a [model of] the traffic and the kind of interactions, the kind of [inaudible] [a handshake]. So ideally, a threat-detection model should contain a combination of log monitoring [IDS]. [IDS] and all the things in a holistic approach.

DARREN KARA: So I think the question is are we looking to do any correlation between the logs being sent back from each of the zones? I think that we could probably do a better job in finding spread analysis or correlation. I don't disagree with that. It's a difficult prospect with the amount of data in servers that we have shipping data back to us.

NOORUL AMEEN: Second part, any meta data on the traffic?

DARREN KARA: Any what analysis? Meta data? I'd have to get back to you on that.

BOBAN KRSIC: Okay, thank you. Thank you, Darren, for your time. We will summarize the outcome and come back if necessary with some questions. And we have a break now for 15 minutes, and see you at 10:00.

STEVE CONTE: We're coming back from break. I'm going to pass to Boban. My understanding is that the subteam is going to be working on drafting some language from the findings that have taken place for the last day and a half, so Boban, I'll leave it to you guys.

BOBAN KRSIC: Oh, thank you. Hi. We will now, until 11:30, use the time to draft the main issues from yesterday's outcome and we'll come back at 11:30

with the next topic that is related to WHOIS compliance and Registry Agreements. So now it's time for drafting and see you all at 11:30.

STEVE CONTE: So Kerry-Ann, we noticed that you just came in the room, the Adobe room. We appreciate it. If you post something in this chat, it does not go to the observer room, correct? So if you can make sure that Kerry-Ann has a link to the Google Doc. I don't know if she's got editing rights. I think there was an issue with her and Google Docs. But then for the record for Kerry-Ann's benefit, if we could keep this unpaused and you guys draft with the intention that she's here and participating.

JAMES GANNON: Kerry-Ann, does that work for you? And if it does, could you write in the chat that that's cool for you? Okay, she's on the flight. So Kerry-Ann, do you mind if we continue to draft without your input and then we'll review back in on the mailing list, or what way would you like to proceed?

Okay, so Kerry-Ann agrees with that for the record.

STEVE CONTE: So then a suggestion, I know that Kerry-Ann ahs challenges with Google Docs, is that once this is finalized today, we turn it to a PDF, at least for a draft, so she can see what's actually written and easier than transcribing it into an e-mail and then going back and forth. Thank you.

All right, welcome back. We've been on pause, so the recording will not reflect that the Review Subgroup Team has been working on drafting notes from the last day and a half. I suspect we'll probably continue that and will be reflected into the wiki once those notes are done.

We are at the slot now with contractual compliance. We have Maguy, and I'm sorry. I don't know. You'll do introductions. Okay. So Boban, I'll pass it to you and/or to Maguy.

MAGUY SERAD: Good morning, everyone. This is Maguy Serad, VP of Contractual Compliance. Joining me in the audience, I have Amanda Rose and I also have Jennifer Scott. We wanted to come prepared.

And we usually anticipate some of the questions you may have for us, but we understand the importance of this exercise and want to be able to provide you and point you to the right information.

To my team members joining me, I know, we want to be mindful of the 30 minutes we have with them. I'll tell you who's who later. But what I want to do, I don't like to presume because the topics on the agenda have changed and evolved. Are you still interested in discussing contractual obligations related to WHOIS as the topics indicate on this agenda? Okay.

All right, well, thank you. A couple of approaches, you have questions or do you want us to kind of give you a high level view of how compliance enforces the obligations for WHOIS? Which approach do you want? Q&A or do you want just a two-minute blurb or how we do it, what we

ask for, where you can find information on data? Because I know Denise is very familiar with us, with the website, with the metrics. But I don't know if the rest of the team is.

BOBAN KRSIC: Well, I propose that we should start with a high level overview, yeah, and then go into a dialog. So it was much easier for us here to get the same level of starting.

MAGUY SERAD: Thank you, Boban. Just a quick briefing here. The first and most important approach in forming this Review Team is that compliance has, and I have prepared, we have prepared a small presentation. We'll be happy to leave it with you and, also, the information is published and Yvette can share the link.

We have created enough information on our website to guide anyone interested in learning more about compliance. We also have a lot of metrics published. But the basic foundation is the complaint types and the approach by Compliance is consistent, whether it's WHOIS, if it's WHOIS inaccuracy format, any other area, unless it's dictated differently like a [inaudible] process might.

Okay, so with that, I'm fortunate to let you, somebody to ask Amanda, our resident expert on WHOIS. And Amanda, if you would share from a Compliance perspective, what do we ask and how do we ensure that the WHOIS inaccuracy obligations are enforced because whether it's WHOIS ARS or it's an external complaint, or it's something we internally

have proactively identified and want to follow-up on, it's the same process. So the source of the information is really less important than the approach and the consistency of enforcing the contract.

AMANDA ROSE: Whether we receive the information through ARS or a community-submitted complaint, if we first make sure we have sufficient evidence to proceed, whether that's evidence of false data, things like bounced e-mails, sometimes it's evident on its face – it might just be blatantly inaccurate – if we do validate the complaint, then we submit the first notice to the registrars. They then have 15 business days to respond.

Should we do the presentation?

UNIDENTIFIED FEMALE: [Inaudible]

AMANDA ROSE: Yeah, that might be good.

UNIDENTIFIED FEMALE: The information is also in that [deck there]. In that link, isn't it?

STEVE CONTE: What we're seeing on the screen is in the chat already.

AMANDA ROSE: Is the presentation on there?

UNIDENTIFIED FEMALE: The presentation is not loaded, but I sent it to you [inaudible].

AMANDA ROSE: Okay, great. Okay, so I'll just keep going. So what we're asking for is that, essentially, compliance with the WHOIS accuracy program specification, that the main things that we're checking are to either confirm that the data has been corrected, or the other option would be that they suspend the domain name if the data is not corrected. They can either suspend, cancel or delete the domain name. There are various ways that they can come into compliance with that. And then the other one is that they can verify that the data is actually correct, which would just be that it's been validated and verified.

So we're asking that the registrars provide evidence if it is either corrected or verified correct, that the underlying RNH e-mail has been verified as functioning. So we're looking for evidence including e-mail headers, correspondence that they sent to the RNH and affirmed the response was received and that they actually did the investigation regarding the alleged inaccuracies.

The other thing we're looking to have them respond affirmatively is that all the data is in the correct format, so there's RFC – I have it somewhere – RFC 5322 would be the e-mail. There's the ITU for telephone numbers.

BOBAN KRSIC: Sorry, can you repeat the RFC?

AMANDA ROSE: The RFC, 5322. And then that's her e-mail, and then ITU-TE164 is the telephone format standards for every country. And then generally, most registrars will do UP validation for the addresses, but they can also use other standard formats for their specific country, so we just look and see that they've actually used a valid format that's not going to be something like putting it into Google Maps and having it come out because that is more of a check and see if the address exists, not necessarily confirming that it's in the correct format.

So those are the main things that we're looking for if the data, again, is either corrected or verified as correct. And then if the domain's suspended, usually, there's several ways that they can do that, either placing it in client hold, deleting the domain name again, or sometimes they'll do it behind the scenes where they put it to a parking page and it'll say something, usually, about the fact that the data is inaccurate, it's [inaudible] RNH go and confirm or correct and they keep it there until the registrant has performed the correct verification, and then at that point, they can go ahead and unsuspend the domain name.

That's generally our compliance process. The original deadline is 15 business days. If a response is not received, then we proceed to second notice. They're then given five business days and, again, if no response is received, it'll proceed to third notice and eventually possibly to breach. But hopefully not [inaudible]. Generally, we're able to get at least a response by that point whether it's one of those three options.

MAGUY SERAD: Thank you, Amanda. To close on our update to you, this presentation, I've made it available to Yvette. She can share it with you. It's a subset of slides that are already available on the link that Yvette also should have shared with you that has much more depth to how we approach compliance, and our goal of that program update – we have one for registrars, one for registries – is to make sure that the community is informed equally on the same platform, what Compliance looks for and how we operate in that space.

In addition to Amanda's update, what we have learned is sometimes whether it's a WHOIS inaccuracy or a WHOIS-related complaint, while the registrar may show that it's been suspended or something, action taken to be in compliance with agreement, sometimes it gets back and it's active.

What the team has done, we conduct spot reviews. We call it quality review and for the domain names that were part of the initial review to bring them into compliance, if they have changed status, the team then sends an inquiry to the registrar. They would notice a changed status, provide us evidence that you've validated or verified, and so on and so forth, and why did the status change. So we try to do a closed loop.

This process is a continuous improvement process. As you all know, the community is very creative and the world is very creative, how they approach any compliance-related matter. As much as we try to anticipate, we continue to learn. I'm not ashamed of saying that, so with

this, we leave the questions to you now and thank you in advance for this opportunity.

BOBAN KRSIC: Thank you. And I would like to hand over to Denise because she has her hand raised.

DENISE MICHEL: Thank you for taking time to come meet with us. I have a whole series of questions. For those you don't get through today, I'll just send them in writing, and of course, if you don't have information or answers, feel free to just follow-up with us in writing.

So which registrars have the highest WHOIS inaccuracy rates currently?

MAGUY SERAD: We do not share this information publicly.

DENISE MICHEL: Does this information require an NDA to get?

MAGUY SERAD: I do not know.

DENISE MICHEL: Could you let us know? I think it would be useful.

MAGUY SERAD: Usually, we go through the [deep, deep] process if there is anything that is related to compliance matters or compliance tickets.

DENISE MICHEL: Right, but we're a Review Team. Some members have signed NDAs. Could you look into that? I think it would be useful to have clarity on what bounds that and the answers.

And so, we received a briefing earlier today about the abuse report that was done for the CCT Review, which indicated there are new gTLDs and registrars where there is greater than 50% abuse of registrations including one registrar where 90% of the domains are reported as abusive.

So I guess my question around this is was that a surprise to you? And has Compliance taken any particular action with respect to this data and these registrars and registries as a result of the abuse report? Thanks.

MAGUY SERAD: Thank you, Denise. As you all know, this seems to be the hot topic in the ICANN industry now regarding DNS abuse. So the team proactively monitors blogs, articles, plus also all types of reports that are publicly available to the world. In addition, Steve yesterday introduced to you a topic that's referred to as DAR.

So Compliance is always on the lookout. I like us to call us like investigators with a little mini [looper], looking for trends, looking for

concerns in the community. I don't want to say if it's a surprise or not because we're proactively monitoring and reviewing different reports and blogs. We do take the initiative to proactively send inquiries when we find articles of this nature that are concerning.

The challenges we face when we send these inquiries is based on either a blog or reports that is all statistical or aggregate and not specific evidence. Regardless, even missing the evidence, we look at it from a compliance perspective, whether it's a registrar or registry, we inquire about it from, "Are you receiving these reports? Are you responding? Do you have abuse contact? Do you respond to law enforcement as required by [318]? If it's a registry, what are your security measures?" So we do proactively initiate those and work to bring into compliance if there is and we find specific evidence of non-compliance, and sometimes if it's a report, we do receive reports from outside, external complaints about abuse with specific evidence which makes it much easier to pursue, factually, an issue.

What I want to also update the team, I'm not sure if you're familiar with the CCT Review Team report and the GAC Communiqué [Annex] 1, they have us to bring more transparency into the compliance complaints related to WHOIS inaccuracy and abuse, and we have been working over the past few months to bring that additional granularity. What that is going to do is in the monthly dashboard, where you see a very generic aggregate number WHOIS inaccuracy broken by the different aspects of it, whether it's an individual report, bulk or WHOIS ARS, we're taking that to the next level to let you, the community know if it's a WHOIS, what type of WHOIS inaccuracy, syntax, operational, identity as defined, and used in the community. We're also taking the abuse complaint and

we'll be providing a breakdown of the abuse complaints, if it's bot-net, malware and the different breakdowns as defined and mentioned in the contract. And we are also working on bringing better granularity to the different areas that were mentioned in the CCT Review and the Public Safety GAC Report. So hopefully, that's going to bring more information to everyone and it's also, we're looking at bringing it broken by legacy and by new gTLD, how many complaints are received in these environments.

DENISE MICHEL: Thank you, Maguy. I guess to make sure I understood what you said, no, it wasn't a surprise that 50% of user registration rates and one registrar with 90% of domains being reported as abusive, so that wasn't a surprise.

And I guess I'm unclear on the second part of the question as to what action Compliance has taken with respect to, particularly the 90% of abusive domains in one registrar, have taken based on that? Or does Compliance feel that it doesn't have the ability to take action based on data of this type. Is that what you're saying? I just want to be clear.

MAGUY SERAD: So I did not say yes it was a surprise or no, it's not a surprise because then I'll be voicing my opinion. It's not a matter of opinion. I'm saying we proactively monitor published reports, we proactively review blogs and when we find and review topics of this nature, we do take a proactive approach absent specific evidence of specific abuse-related issues, we do initiate what we call an inquiry. An inquiry in the

compliance process, it's like a fact-finding approach because we don't have specific data. We do ask the registrar or the registry, specifically tying it to the contractual obligations on that. So we do follow-up and the action will be depending on the responses and it's all on a case by case scenario.

STEVE CONTE: You mentioned the data. I know that we just swiped our dashboard out and we're now doing accountability indicators. Is the information you mentioned on the data board, is that in the new accountability indicator dashboard that you had mentioned [would reflect]? I'm sorry.

MAGUY SERAD: Yeah, the dashboard I'm referring to is on the Compliance page on ICANN.org. If you go on the Compliance page, there is performance measurement and reporting. Compliance publishes monthly dashboard on the Compliance operations and activities related to complaints and the processing of complaints and resolution. That's the monthly dashboard that we are enhancing to bring more granularity to different types of complaints.

STEVE CONTE: Thank you. I'm just trying to find the second posted for the room.

MAGUY SERAD: I will send you the link.

STEVE CONTE: Oh, thank you.

DENISE MICHEL: So does compliance treat registrars with higher inaccuracy rates to greater scrutiny and compliance activities?

MAGUY SERAD: I'm sorry, Denise. Can you repeat the question?

DENISE MICHEL: So where we see registrars with higher WHOIS inaccuracy rates, do you treat those registrars to greater scrutiny? Do you increase compliance activity on those?

MAGUY SERAD: So we do look at the data internally from an operational perspective to see if there is a high rate of a specific type of complaint issue, whether it's WHOIS inaccuracy or another area. And what we do in that space, we start to look at it. While we want to address specific issues reported to us, we also take what I call a bigger systemic approach to the issue to try to understand 1) the first thing is let's address the reports that have been filed with us, 2) what are you going to do differently to stop this behavior and fix it, and what we call that is what we ask the registrar or the registry to provide us a remediation plan.

They have to come forward with that solution. We do not tell them what to do. We don't tell contracted party. They own the remediation plan based on the issue and the case we're dealing with. A remediation plan is reviewed by Compliance and followed through.

Now, in the future when we observe the same behavior, we don't go through the one, two, three process because they have already used that and what we do is we take it immediately to a third notice, which we refer to as an escalated notice and we ask them, "Why is this happening? Correct the issue and let us know what else you can do about it."

As we all know, Denise, you may have worked with Denise long enough and I'm trying to anticipate her questions. Let's see if I'm going to be successful. As we all know, WHOIS inaccuracy data, some of it is recent, some of it is grandfathered, so yes, while the registrar may have corrected because they went on 2013 RAA and are doing different approaches, it does not mean that all data has all come up to speed to it. So it's a balance where we are trying to kind of work and bring everything into compliance based on the cases brought to us.

NORM RITCHIE: Are all registrars on the 2013 agreement now?

JENNIFER SCOTT: This is Jennifer Scott from Compliance. No, not all registrars around the 2013 RAA, there's about 20 that are left on the 2009 RAA version, and

of those 20 or so, some of them don't have names. So the bulk of the names are being covered by the 2013 RAA.

DENISE MICHEL: So what trends is Compliance seeing in the [periodic] reporting, that is are WHOIS accuracy rates going up or down? What's the trend data there?

MAGUY SERAD: Can you link to the performance report? There are a couple of reports. Accuracy rates are always the higher volume. WHOIS inaccuracy is the highest volume of complaints we receive. But we also have a certain months trending report that can show the volume of complaints by complaint type over the past 13 months, if that's valuable. I need to refer to the metrics. I don't have it memorized. Sorry, Denise.

All right, I am looking. If you have opened the link to the performance measurement and reporting, the 13 month trend report will be the one titled "Informal Complaints by Region". If you go to the bottom of it, you will see the trend from August 2016 to August 2017 and I'm going down to WHOIS inaccuracy.

What I want to point out is while the volume for WHOIS inaccuracy may seem higher at times or even on the monthly dashboard, you can see the same, please keep in mind that some of this data is also driven by the WHOIS ARS exercise. The best way to see that, but it doesn't show it here, by complaint type. But if you look at our monthly dashboard in addition to that informal report trend for 13 months, you will see, for

example, in August, not in August, in September, in December, February, these are the areas where you see a higher trend in volume of complaints. If you open the dashboard, you will see a WHOIS ARS reporting activities that took place and made it into the compliance process. And that data is all available. I don't have this information memorized, but it's also available on our monthly dashboard, and on our 13 month if you're interested in following through on it.

DENISE MICHEL: I'm sorry. I don't have access to that link, so do you have a sense of whether accuracy rates are going up or down broadly?

MAGUY SERAD: I'm not sure how to answer that question, Denise. They fluctuate. We measure it on a month by month and we look at the trends over the 13 months. Like I said, they go up when we have the WHOIS ARS feeds come into it, but from external complaint received before we provide it here in the monthly dashboard – let me see, just a second – so they fluctuate. It's not like it's always on the [grove], or it's not always on the decline. They go up and down based on what's the source, how much activity is coming in from the WHOIS accuracy reporting or from bulk submissions.

BOBAN KRSIC: So we have only two minutes left, and I would say thank you before we stop this topic. And I have a few questions off from the WHOIS accuracy and so on. What about other entities and your responsible-ness to

them, like the bureaus or data escrow providers because they are also mentioned in the RAA and you're also responsible for them. Have you ever something provided to them or how you measure this requirement to them or can you give us a little bit more information about entities like DACAs or the bureaus and contractual complaints?

MAGUY SERAD: Good thing the recording doesn't capture our facial expressions. I'm sorry. I'm staring at Boban, trying to understand. Are you asking how does Compliance enforce contractual obligations with data escrow agents, the EBERO vendors?

JAMES GANNON: [Inaudible]

MAGUY SERAD: I'm sorry?

JAMES GANNON: [Inaudible]

MAGUY SERAD: Monitor and, okay. And that was James. I was repeating what James said.

Okay, so from a contractual compliance obligation, as you know, our contracts are directly with the registrars and the registries to enforce the obligations. Let's just simply look at the data escrow.

As you know, ICANN subsidizes data escrow services via Iron Mountain, but also, there are other data escrow providers that have been reviewed and approved if registrars want to engage with from a registrar perspective, from a registry to different scenario.

The obligation in the contract is for the registrar and/or the registry to deposit based on the frequency and the format and the specifications. So we work directly with the registrar and the registry on that space. We proactively also audit when we are conducting audits to make sure they are fulfilling their obligations in that space and it is their responsibility to work directly with a data escrow agent in that space.

For EBERO, EBERO is a parallel process as you heard yesterday. I have been one of the observers. I am listening to every question you guys are posing. Yes, for the EBERO, as you heard from my counterpart, it's a different process that is led outside of compliance activities, and James, you submitted a question on that based on what is the compliance activity in that space?

As you heard Francisco describe yesterday, there is a strong handshake when there is a failure in that obligation to a specific threshold and please don't ask me to quote those. I don't have them memorized. That triggers a compliance notification and based on the severity of that notification, while our goal, it is to bring a contracted party into compliance, the first and most immediate goal is to make sure that we

are securing and stabilizing the Internet. So that's why Francisco's team takes priority in that space over compliance. Parallel to that, the process, we do reach out to the registry provider to figure out what have you done, what were the remediations in place, and how are you going to be in compliance? And we see it through completion. We do depend on technical services to help us asses and review that they are in compliance because it's a different technical skillset. So it depends on the cases we're dealing with whether it's proactively through audits or monitoring, or via external complaints.

BOBAN KRSIC: Thank you. [inaudible] actually designated escrow agents and the accredited escrow agents, they are the third-party providers. And the revocations. Do you also monitor them and report it? Because you told us only the registrant, registry side, not the escrow provider side. What about that?

MAGUY SERAD: I'll ask Jennifer to speak more to it because there are some notification obligations in that process. Jennifer, do you want to elaborate on that, please?

JENNIFER SCOTT: Sure. So if we're talking about the registry data escrow space, and specifically the EBERO thresholds, there are requirements for ICANN to notify the data escrow agent of the impending failure. But I think your question might be a little bit different, you're asking if we monitor and

make sure that the third-party providers are staying true to their obligations under the contracts either with ICANN or with the contracted parties.

We do occasionally get complaints from contracted parties that they're not feeling like they're getting the service that they would otherwise want from the third-party providers, and so where we have a stronger handshake, like with Iron Mountain on the registrar side, we are able to reach out to them and liaise with them to address those concerns.

In terms of the other DEAs that are approved by ICANN but are not subsidized, it's a little bit less of a stronger handshake, but we still do have the ability to reach out to them and discuss anything that the contracted parties might be bringing to our attention.

JAMES GANNON: One minute, don't worry. Coming back to the question that I've actually put in writing, and I'll ask this simple question. Within the compliance framework and when you're trying to enforce contractual obligation, there is a decision point when you move from working with your contracted party to when you move to enforcement actions.

Is that formalized? And if so, at what point do you have decision points on escalating these complaints from, "Okay, we've identified an issue, we've reached out to a registry and they're not responsive?" Is there a formalized framework that you work within for those levels of escalation? And at the end of the day, at what points do you move from working with to enforcement and rescinding the contract, essentially?

MAGUY SERAD: Thank you, James. This is a topic near and dear to my heart. I'm all about processes. Yes, it's also on the website under the approach and the methodology. It's in the link I shared. I'm happy to point you to it during the break.

The contract doesn't speak of an informal and formal resolution, just says you're in breach of your obligations. We created an environment over six and a half years in collaboration with the community to say, "You have got to allow us the ability to work informally and address issues," because not every issue is really – sometimes it could be a language barrier, it could be a misunderstanding, or it could be just us misinterpreting information or we have wrong data.

So we wanted to have an informal approach. We call it informal resolution process one through three. If we do not have facts, as you heard me earlier, we call it an inquiry. If we have facts, we call it a notice.

Through that process, it's a collaboration process, but it's also a structured approach based on the complaint type. Absent a contractual obligation in the complaint type, like the WHOIS inaccuracy requires 15 days, we created one that says 5-5-5. Five business days, and sometimes we may do a follow-up because of this one little thing missing, or there's collaboration but it's missing a few elements.

But if there is no collaboration at all, it progresses to the second notice, third notice. By the time we reach a third notice, then that's the last and final. We issue a notice of breach. Everything in the enforcement phase

is published on ICANN.org forever and ever. And we update the status of that notice of breach based on how it's progressing.

The simplest status is notice of breach issued on date X, and we provide a chronology and a summary of the issues. Then the simplest status would be updating it with a cure date. Sometimes the cure doesn't happen. We may have to suspend or extend or follow up, but that status is continuously updated until it's either resolved or what we call cured, or escalated to the next level of enforcement.

Did I answer your question? So there is a structured approach for enforcing, and it's all documented and reported on.

JAMES GANNON: I'll write a follow-up question, I think it's easier.

BOBAN KRSIC: [inaudible] time I would say. Thank you all for being here, for sharing your time. There are a lot of questions, and I thank you. And we will summarize it, and if there are some questions, we will send it out, and then we'll be happy if you can answer it. Thanks all, and see you at 1:00 p.m. with the following up stuff and security incident management is waiting at 1:00. Thanks.

UNIDENTIFIED FEMALE: Thanks, everyone. Just before we get started on lunch, Theresa and Larisa wanted to stop by and just briefly chat with the team.

UNIDENTIFIED FEMALE: I won't interrupt for too long. I know you guys have been incredibly busy and had a really productive time, at least that's what I'm hearing. So welcome to L.A., and I hope the weather is a little conducive to everything. And if there's anything you need or any other support from the team, just let me know and we'll make sure that happens. So I hope it's been useful, and if you need anything, I'm out and around in the halls. It's good to see some of you. And is it going well?

UNIDENTIFIED MALE: Yes.

UNIDENTIFIED FEMALE: Yes? Okay. Great. Well, thanks. I don't want to keep you from lunch.

STEVE CONTE: Thanks, all. And all the links that Maguy and her team have referenced, including RFCs, I've put into the Adobe room chat. So we have it for the record for review upon later reflection.

UNIDENTIFIED FEMALE: Yes, Boban. Thank you. Apologies for the delay, everyone. Gary unfortunately has been part of the evacuation in Orange County due to the fires, so he's not in the office, unfortunately not able to join us. But Sam, the other SME on the list, is able to join, and here he is right now. So we can get started with the conversation. And of course Sam only

knows part of the discussion, so we'll see how much help he's able to provide to the Review Team today, and we'll take it from there.

STEVE CONTE: So we're recording, we're back on. So please when you speak, use the mic. I'd like to welcome everyone Sam Suh, he's our VP of Back Office Solutions. This session is the security incident management portion of the agenda, and I'll leave it to Sam or Boban on how you guys want to start this session out. We have 45 minutes. No. an hour and 15 minutes, and then we have some continuation on this if needed after a break. So right now, we'll go until 2:15, so we have about an hour right now for discussions. So Boban, I'll give it to you.

BOBAN KRSIC: Thanks, Steve. Hi, Sam. Welcome. Thank you for your time. We'll talk in the next hour about security incident management at ICANN and how it is related to DNS SSR. It'll be good if you can give us some overview of the process you have in place, and then we'll go into a dialog here to pick up specific points and discuss them.

SAMUEL SUH: Sure. Hello, everyone. Thank you for having me here. Let me just go into my role a little bit, and then we'll go into the security incident management process.

I'm Samuel Suh, I'm part of the IT team, and I work for Ashwin Rangan. Unfortunately, he couldn't be here today. He's dealing with the fires in Anaheim.

If you want to look at IT in the two homes, one would be community-facing, supporting the community. The other would be internal facing. I'm representing the internal facing side. Some of the process that we'll be going through actually covers that, because the majority of the incidents, if it occurs, obviously it hits the outside. But all of our internal processes happen here.

So as far as incidents are concerned, as we get notified or as we uncover the incidents, there is a workflow that we follow. First, if an incident occurs – let's take a website or something like that – we have a website that goes down.

We have a support model where somebody or a group of individuals are on call. So those individuals that are on call will be notified through our systems if any systems or processes are out.

That individual will try to triage that as best as they can. If something happens beyond their skillset, then there's a pool of other resources that are on call as well.

So for the most part, the majority of our IT operations as well as the executives, we all have – there's a chain of command in which it goes for incidents, so it goes up the line.

Once the incident has been identified, if we can resolve it, then we will resolve it right there on the spot. If we cannot, then we take it to the next level and see what we do.

The biggest thing is recently we did have an outage in one of our websites, so I'll go through that example. Fortunately, it was an update,

a patch that we were able to identify and we were able to roll that back to make sure the website was still functioning.

Subsequent to that, there was an incident that [inaudible] down. There's a short-term and a long-term fix. Obviously, the short-term fix was to get the website back up. The long-term fix is that we look at our processes of how the patch process works.

After that, we'll go through a security review again, and there's some redesign that we need to do for this particular incident where we need to make sure that we have multiple failovers versus a single failover. So things like that. We're always in a mode of continuous improvement and trying to make sure that we're meeting our service levels [inaudible] time.

But every incident – and it depends on the level of severity. So if it reaches a critical severity, we're disclosing. We disclose everything on ICANN.org. If we have an incident, an outage or what have you, that's in our SLAs. With our indicators page, we show our uptime and downtime, as well as if we have a major incident, we report on the root cause and the fix, and any other subsequent relevant information. Yes.

DENISE MICHEL: Sorry to interrupt. Where's this on the website?

SAMUEL SUH: Right now, it's kind of scattered. You have to search. But I think within the next couple of weeks, we'll be reporting on all incidents on a single page.

DENISE MICHEL: When you have a chance later, could you e-mail us the links to the various places that it's currently located? That would be useful. Thanks.

SAMUEL SUH: Yes. I have the last two. I think we have the incident – there was a policy that was written way back when that's still valid. I think we're in the process of updating that, and then there are two incidents that we have that have been scattered that I can send to you guys.

So that's the basic process. the majority of those incidents we all track, and through our internal systems we make sure that each of them has an incident report, the fix, and then depending on the severity, disclosure.

NOORUL AMEEN: You were mentioning two incidents. One I can see that that's a possible spear phishing campaign targeting ICANN organization 2014. And you have an update on 21st February 2017 regarding the possible disclosure in .net or something like that.

Can you explain the incident process for handling that [inaudible] breach?

SAMUEL SUH: You're talking about the Salesforce? Can you repeat that?

NOORUL AMEEN: The spear phishing attack in 2014.

SAMUEL SUH: I'm not sure exactly what happened with that. The procedure is the same. If we have a spear phishing incident, it's exactly the same. So on that one, it would be our Security Team which is headed by Geoff Bickers. They're notified. We have an anti- spear phishing campaign as well internally that we actively utilize.

But the process is the same. Once it's been identified, we try to prevent or stop what's happening, and then we go through a process of [cleaning,] the entire process and then recording the event.

NOORUL AMEEN: Two questions. First thing is as a proactive measures for this kind of [inaudible] campaigns or spear phishing attacks, organizations have their own solutions to detect in a proactive manner before reaching to the ultimate customer. That is the proactive detection part.

Second thing is the reactive part. In the spear phishing campaign, you could have done some analysis. So that containment might have been happened, I think. So, could you be able to pinpoint what is the possible source of attack? Or I don't want to name a particular person or country, but what might have been the reason or the root cause analysis kind of thing? Actually, that is a part of incident handling.

SAMUEL SUH: Are you familiar with that?

STEVE CONTE: I'm somewhat familiar. I was here, so I guess I'm one up on you on that. Not speaking of the phishing attack, I can speak on some of the other types that we've seen back when I was involved with IT, and some of the causes of that have been everything from visibility, we're a visible organization and a tempting target on various types of attacks on that.

Some have been a disruption attack or a disruption attempt when there have been particular issues, community issues or public issues that had a divergence of community opinion. So some new TLDs in the past – not any of the new group that I know of – we had TLD applications and groups were strongly opposed to that, so we had – I would call it a very active activist campaign rather than an attack, but it was meant to take notice and disrupt systems.

And root servers are always a juicy target, so there's the whole, "Put a feather in your cap, I've attacked a root server or I brought a root server down." That's largely changed because of the Anycast model. It's hard to bring – not impossible – the root system down. There's a myriad of ways.

Part of it is also just pure opportunity, and we're not the only organization that's subject to attacks, especially with spear phishing and all that. And I think in some ways, we have to look just from a regular, organizational perspective, people are going to try to get into other organizations and see what they can find once they get in. That's part of the nature of hacking and attacking.

So it might have been just pure curiosity, it might have been trying to get a leg up and trying to get that next step in to penetration. Without having the data that IT might have or the analysis that might have taken place after that, I can only speculate on why we were the target of the spear phish, but I wasn't surprised that it eventually happened, and I think after the mitigation of the attack itself, I think the actions of IT, Geoff Bickers and the Security Team were entirely justified and valid.

And because part of the spear phishing is a social attack, and Sam just brought up the anti- spear phishing endeavor that took place, I think that was really important because as an education process that if it's a social attack, the weakest link is the human, and we have to address that. So that's a very active campaign within the organization since that.

I know that wasn't a complete answer to what you were asking, but I think it's hopefully enough that it addressed your question.

NOORUL AMEEN: There are APT solutions which will prevent these kind of malicious attachments and those kinds of issues. So, are those kinds of systems in place?

SAMUEL SUH: From an e-mail perspective, we go through a third party, and they do the active monitoring and try to stop those. If it does come through, we also have very different levels of security. Let me just go through I guess from an internal perspective.

There are some things that we've actively done to increase our security internally. One is educate the person, because that's the weakest point. Every year, the entire ICANN staff has to go through a certification process where they have to look at all the security videos. We have a Q&A session afterwards. That's an annual event that we have now. I think it's following that 2014 incident.

We also have an active spear phishing campaign where we actually promote false spears in our e-mail system and see how we do, and educate the staff on what to do and what not to do.

The third thing is our password, our firewall security has been increased by many folds, and our two-factor is stronger than it's ever been. We have a true two-factor authentication. We have very complex passwords that we have to use now. We have to go through a VPN tunnel for anybody who's remote.

So those are all the things that we've done. And through our e-mail – we're an exchange house and we go through a third party, and there is an active participation of things that they stop and block based on the filtering rules that they've set up that we've agreed to.

NOORUL AMEEN: That is particularly regarding the spear phishing attack. As an organization perspective, you may be seeing different kinds of incidents, like in other attacks – phishing, scanning, malware and spamming kind of incidents. What is the percentage of these incidents on an annual basis or something like that?

Second thing is that how you're doing the root cause analysis, the log analysis or mitigation parts and all those kinds of issues. What is the –

SAMUEL SUH: I don't have those data points, but depending – I guess we could get that for you as to what we've – there's our security records –

NOORUL AMEEN: [inaudible] generic trend you can tell me [inaudible]

SAMUEL SUH: It's very low.

STEVE CONTE: I think even reporting back – and unless we're talking about security logs here, I want to be very careful on what we're saying we'll get back to you on. Unless we have a very specific ask on what you're looking for on that, I think generalizing, trending is going to be the answer on this without – and this falls into that Category 5, exposing ICANN's risk and not wanting to. And so I just want to be very careful on what we're going to provide as far as logs and data from those logs go.

SAMUEL SUH: Yes. It could absolutely be used as a recipe, so yes.

NOORUL AMEEN: I'm not asking about [specific] questions. Normally in a website [inaudible] website intrusion, your web server logs or perimeter device logs will tell you the kind of intrusion attempts and [inaudible] Maybe some kind of – if it's signature-based attack, it'll get some kind of detection from the log analysis.

So my question is whether the process is there or not. I don't want to see the exact logs or what kinds of logs you're maintaining. And log policy [I can ask, except for] how long you maintain the logs and what kinds of logs you're maintaining and what is the detection process and how the root cause analysis doing.

So I don't want the logs, [inaudible] things, but a generic perspective of how the root cause analysis is doing, and these kinds of issues.

SAMUEL SUH: The answer is yes, we do it. It goes through the same exact process. We log everything. There's probably nothing that we actively destroy, so we have the data points to go back to if we need to recreate. We have backups of everything. But every single incident is logged, and we track it. I don't know what else to say.

JAMES GANNON: One question, and then probably some follow-ups. Do you differentiate between, let's say, IT incidents and security incidents, or do you manage them in the same way?

SAMUEL SUH: Depending on the severity of the security incident, the reporting protocol is different. But the process itself, if it's an incident, it's managed the same way.

JAMES GANNON: And then my follow-up, again I suppose it's a differentiation question. There are what I would call business IT systems which are to run the ICANN corporation which we're not terribly interested in, and then there are the more critical DNS supporting, root zone supporting –

UNIDENTIFIED MALE: [inaudible]

JAMES GANNON: Exactly, yes. So again the same question, is there a specific set of incident management processes around those with different escalation points to the standard [what's called] business IT systems?

SAMUEL SUH: The process is again pretty much the same. The escalation is different in that there are different resources that support, but the process is consistent. And that's the best way to do it – for me, because one individual could be supporting both sides of the house, as an example. They should be following the exact same protocol.

NORM RITCHIE: Is there any separation between the systems supporting IANA functions and the rest of ICANN, or is it all the same network and systems?

SAMUEL SUH: Same network, same systems, other than the L-root's all over the place.

NORM RITCHIE: Okay, so in that regard then you could have traversal, you enter at one point and traverse over to IANA functions [as part of an] intrusion.

STEVE CONTE: I'm stuck on your word "IANA function." There are different aspects of, as you know, the IANA function. So yes, I guess in some ways I would agree that if it's a systematic function of IANA, then a compromise might expose that. But much of the IANA function beyond root zone management is just documentation and being a repository for unique identifiers for IETF and other bodies. It goes back to if you could narrow down a little bit on your questioning.

NORM RITCHIE: Yes. I guess it would be whatever machines – I assume it's machines are used to manage the root zone, or the numbering or port assignments.

ERIC OSTERWEIL: Basically, the registries that manages for IETF.

STEVE CONTE: I would feel more comfortable if we took that question and discussed it, and then provided some sort of answer, just because I want to make sure that we're being accurate, but also, this falls on the point of what can be publicly said and what can't.

ERIC OSTERWEIL: I just [spilled something]. Yes, I think that's great, sort of a good, well-reasoned answer I think is what we're all looking for. But just to sort of state at least the concern that I think I share with some of us – and not to assume I know what everyone is talking about, yes, the extent to which there's breachability or dependencies that exist between systems and repositories that support either IANA's operations, its systems, its registries or anything else with other things that are managed in the corp IT infrastructure, to the extent to which we'd want to know, is there any sort of compartmentalization, are there any boundaries? What sort of protections are in place to isolate various pieces so our printer is not able to access because they have to go a segmented [inaudible] or something like that? What are the controls that are in place to protect a beachhead from reaching what services the global identifier system in some way?

NORM RITCHIE: You're absolutely right, that should not be responded to here.

STEVE CONTE: Unless you want to comment right now, I completely agree that some of that question is logical, and I don't want to answer or even ask for an

answer right now until we have an idea of where that threshold of information sharing is. Thank you.

NOORUL AMEEN: I just want to know more details about your security operation center and the designated role of [CSO.] What kind of functionalities the [CSO] will look after? And do you have any officers specifically designated for the [CSO]?

SAMUEL SUH: One of the [legs] reporting in to Ashwin is Geoff Bickers. Geoff Bickers acts as our [CSO]. The role, he has a team of resources, and they're responsible for basically the security of – so security standards, security protocol, the minimum thresholds when we vet a potential internal tool, as well as any websites that we go through.

So it's a combination of I guess things that happen. First, if we're doing some custom development, it stats the development team to make sure that they're following security standards. Then it goes through a two-way process where they have their security standards.

It also goes through a Security Team where they put together their security assessment for an application or a service whether it's internal or external. But that's the group that sets the rules, the boundaries, when they should be – as part of any active implementation or campaign that we're working on, they're part of the team.

It's one of Geoff's representatives is associated being a part of the team representing the security for the group. So it's a fully functioning group with several heads in it.

STEVE CONTE: Can I interrupt for a totally unrelated item? I'm being told we're having a fire drill today at 3:00 p.m. They need to put this room out. It's a standard drill, it happens every once in a while. It's common practice. We're going to have a fire warden – whoever that is – come and get us at 3:00 p.m., so we do have to evacuate the building, just until the bells go off and stuff like that. I'm sorry to interrupt you, but I'll forget if I don't say it right now. So 3:00 p.m., we're going to plan a little break. We're going to go outside, we're going to enjoy the fresh air. And now I'll return you back to your regular program.

NOORUL AMEEN: One more question. Basically if you consider incident handling, there are two kinds of processes. One is internal incident handling process, and the second thing is that global issues [and all things.]

I will give you an example of this. If you have a squatted domain similar to ICANN.org or .com, the incident handling procedure would be different because the entities involved are different groups. And if you have a targeted attack or a scanning attack on your website, that is your internal issue, because all the access and logs are managed by [inaudible]

So you have any policy to distinguish between these two kinds of incidents [inaudible] the same team or the same kind of process you have following for both incidents?

SAMUEL SUH: That's a good question. I don't know the answer to that, so I think that's one of the items I'm going to have to get back to you. I don't know if Geoff has two sets of policies, one external and internal.

STEVE CONTE: And Sam, just to clarify, we are recording this session, so we will collect the asks that are relevant to you and your team, and we'll pass them on to you.

JAMES GANNON: First of all, I'll ask, does the Salesforce deployment fall under your remit?

SAMUEL SUH: It does not, but I'm aware of it.

JAMES GANNON: Okay. I'll ask questions, and if you can answer them, great. if not, we'll take them on the record. My understanding is Salesforce – and potentially other cloud-based solutions – are used as part of some of ICANN's operations, which at times are used in DNS supporting roles,

particularly during the new gTLD application process and in some other – yes, I can see Steve having a confused face.

So to the best of your knowledge, could you walk me through the vendor management and security management of a third party vendor that ICANN does particularly with Salesforce? Because I understand that it's a pretty big deployment, but equally generically to any other particularly cloud-based providers that you might be using in more critical areas. Again, those business IT stuff that we're not terribly concerned with in the more impactful areas.

SAMUEL SUH: I'll answer a little bit more generic because I don't want to name the groups. But for Salesforce, yes, it's the Force platform. So what we've done is once we set up an environment or developed something in that we've actually asked Salesforce to come in and vet our solution from a security standpoint. That's the first thing that we've done.

The second thing is through our Security Team, there's an outfit or a group that we bring in to assess the security standards. So it goes through two pretty detailed layers of security, one with Salesforce directly. The second thing is with a third party independent assessment as well, so that's the level of scrutiny and security that we go through.

JAMES GANNON: Is that process – likely in a more [lightweight] manner – renewed on, say, an annual or a bi-annual basis to ensure continued compliance rather than just the initial implementation?

SAMUEL SUH: The security assessment is continuous. The third party assessment, I don't know.

NOORUL AMEEN: SSR role, we do both proactive and reactive measures. As a proactive measure, we publish [inaudible] notes, advisory, security guidelines and wide SLRs and all the things. Secondly, as a reactive part, we are handling incidents and doing the root cause analysis on all the things.

So I've seen some alerts regarding the WannaCry infection and all the things on the ICANN website. So ICANN is addressing many issues as a proactive measure, because ICANN is not in the operation part mainly. So whether it's a part of you Security Team [inaudible] separate team is for there for publishing these kind of advisories and [inaudible].

Second thing is that, do you have a cybersecurity policy or IT security policy and a crisis management policy kind of thing?

SAMUEL SUH: Yes, we do.

NOORUL AMEEN: First part, proactive and reactive?

SAMUEL SUH: The proactive, from a policy standpoint proactive, I don't know if there's a set standard other than the security things, the training that we have to do. Whether that's in a policy or not, [inaudible]

NOORUL AMEEN: Some team is involved in the preparation of the advisories and alerts, [just I've] seen in the WannaCry or Petya incidents. Is it a part of security operations group, or [inaudible] this is actually happening or a separate research and development team is looking after?

SAMUEL SUH: It's the same.

JAMES GANNON: Another question. And I suppose it builds on Ameen's points. Do you do active vulnerability scanning within your internal network to identify known risks within your existing application space? And if yes, who owns those vulnerabilities for remediation, and who owns the action plans to bring them back into compliance when [inaudible]?

SAMUEL SUH: The answer is yes. It depends on where it's found. We have different IT heads associated with it. For me, we use Oracle for our ERP system. If there's a known vulnerability with Oracle, it's a cloud-based system, I'm actively chartered to work with Oracle to make sure that it's resolved.

As an example, the recent incident that we had with Equifax that you guys may be aware of. Apache [threads] was kind of the root cause that came up. What we've done is go to Oracle and ask, say, "Hey, look, we know that this was something that was released. We know and assess and understand the vulnerability with our ERP platform. Give us a green light, a red light or a yellow light as to the various applications."

It's not an easy ask. Oracle is a gigantic company. But it's something that we pay for, and so we're in the process of collecting the data for the various systems that we use on the Oracle side. The answer is pretty much – there's only one application we're waiting on, the answer is no, that they went through and vetted everything. So that's the process that we follow. It depends on who it lands on. And if it lands on ERP as an example, that'll be my responsibility.

UNIDENTIFIED MALE: Just a clarification on the line of question that could take place. Oracle and ERP is some of our internal systems that we use for HR, accounting and things like that, and would be up to you guys to determine if it's SSR-related or not. Thank you.

NOORUL AMEEN: This is regarding the vulnerability reporting things. As I said, ICANN.org is using Apache and Ruby on various kinds of platforms. A couple of times, as a general user [inaudible] finding out a vulnerability that you may be using a vulnerable version of Apache or something like that. How to report this thing? Do you have a mechanism to report this

vulnerability? And how are you addressing that vulnerability, and how are you fixing those kinds of issues?

SAMUEL SUH: There are many ways. I think there's a feedback section on probably any one of the pages. The other one is our ARR process where you can go ahead and log a question.

STEVE CONTE: I can speak to that if you want, but I think we could do that out of session. It's a process through Salesforce in which we receive and track communications to Board, communications to the CEO, advisories from the SSAC and RSSAC. This way, we can track the process to make sure that they're not falling in holes and not being responded to. But I can go into further detail at a later point when we're not wasting Sam's time on that if you want.

BOBAN KRSIC: Yes. Thanks. The security incident management has an interface to the global risk management. So if you had an incident and you have a fix, and you have lessons learned, provide lessons learned, what's the outcome of them? Is there an input to the risk management to address such incidents in appropriate way when they come again?

SAMUEL SUH: I can see the tie. Yes, I think the best way to respond to that is we're in continuous alert. We are very paranoid about things. There things that

we have to do that I've not done anywhere else. I've been in the IT field for many years, and the level of scrutiny, the level of review is very rigorous for us to stand up a platform. Just to log in in the morning is very difficult here.

So yes, everything is tied. Whether it's a risk assessment, whether it's looking at our – we're in the process of replacing ICANN.org. it's a constant state of review and upgrade. And a little bit of paranoia.

NOORUL AMEEN: I feel there's a confusion, because the same persons are involved in the operation part and the same persons are involved in the security part also. It's a conflict of interest, I feel.

SAMUEL SUH: No, we have a separate team for the security. The security operation, when it comes to maintaining our internal or our community servers, infrastructure, that team reports to me. Geoff Bickers and that group reports – there's a security leg that's the security guys.

NOORUL AMEEN: Okay. And normally in a critical sector organization, they will outsource the security services, for example the banking sector. According to [inaudible] standards, there'll be QSAs for doing the security services.

So ICANN is having an internal process for a security service, or some of the parts were outsourced to some other parties [inaudible]

JAMES GANNON: Sorry, I'm going to come in. I slightly disagree with your interpretation of what a QSA does. I used to be a QSA.

UNIDENTIFIED MALE: [inaudible]

JAMES GANNON: Yes. No, just – I can elaborate on what I think the intent of the question was. ICANN appears to have a unified incident management process which involves both the operations and the Security Teams converging into similar roles, if that's correct. I'd be interested in knowing where the handoff between what – let's call it an IT incident becomes a larger security incident which is then handled by Geoff's team. Where does it morph from being a vulnerability that you need to patch as IT into something that needs security incident response, which is very different? I have concerns about it being under one process, but in my mind there are sometimes different processes that need to be followed.

SAMUEL SUH: The operations folks are responsible for fixing. If something is identified as a security issue – and when we know once we identify it, the process of fixing wouldn't be with our Security Team. That's with our infrastructure group. They're responsible for making sure that the Band-Aid is applied and the long-term fix is applied.

What security folks do is they make sure that we're following protocol and making sure that the same vulnerability will not continue, as well as handle more of the reporting aspect saying, "This is the incident," following that chain of command, working with the Legal Team or the Comms Team to formalize an incident and post it and let everybody know. But that management process is with our Security Team. And the rules associated with the guidelines, rules, and whether it's officially been patched.

We can say that we patched something and we believe that we've corrected it. We'll go through our own internal process of validating that, but it's really up to the security team whether they decide to do things internally or decide to do it externally. Depending on the skillsets that they have, it's all up to them. They act independent.

ZARKO KECIC: Do you, and how often do you do [inaudible] of implementation of security policies?

SAMUEL SUH: At minimum, annual. But we're going through a process right now of validating that. So we've gone through, let's say in the last two years plus, almost every six months we've gone through some sort of a third party assessment on just nature of IT, the nature of our network and security of our network, review of the internal practices that we have.

So it's been probably – I'd say within the last two years, it's been more frequent, only because we're in the process of establishing policy and

procedures, setting policy and procedures, and we've actively gone out to third party to vet and ask questions. So it's been more active.

ZARKO KECIC: How about internal checks?

SAMUEL SUH: Like quality assurance?

ZARKO KECIC: Yes.

SAMUEL SUH: We have a party of one on our two-way side. He's recently joined, so we're actually working with him right now. The biggest issues and things we're working on is the business continuity plan, as well as the disaster recovery plan and execution. Disaster recovery, we exercise that pretty regularly. But formally, every six months.

JAMES GANNON: Question for interest, and not sure if it's respondable or not. What would the split within back office system be between custom developed and commercial, off the shelf? Because they have a different level of risk. Just in broad percentages.

STEVE CONTE: Let me ask a question before your question. What back office systems – and fix this question if you want. What back office systems do we feel arguably help manage or impact the unique identifiers that ICANN is responsible for? And out of those, which are – to James's question – outsourced and which are inhouse? Does that fit within your scope of questioning? And –

SAMUEL SUH: I can give it a broad answer. I would say when nit's more internal-facing, you'll find more off-the-shelf solutions. Systems that we use internal you'll tend to find things a little bit more custom [inaudible] or developed inhouse, engineered I guess. But I don't have a percentage.

STEVE CONTE: And that's based on community-driven demand on what they want, such as the Centralized Zone Data system or whatever it's called.

JAMES GANNON: And then my follow-up was going to be, is there a different level of rigor applied to the management and auditing of your COTS systems versus your custom-developed?

SAMUEL SUH: Yes. One somebody else developed and got it to the point of industrializing it and making it available to the public, the other one is something that we've engineered ourselves, so it's way different. Totally different lifecycle for some of these things.

NOORUL AMEEN: This is regarding the reactive handling of a particular incident. Suppose you're observing any kind of probes or attack [inaudible] from a particular country IP, so for country X. [inaudible] handling procedure is to contact IP owner or the domain owner of that IP and you need to report it to the corresponding CERT or the country CERT. Do you also follow the same kind of procedure, and do you report incident to concerned CERTs also?

SAMUEL SUH: I don't know the answer to that.

STEVE CONTE: I think participation in CERTS is I feel a valid question, and if it's alright – I'm getting nods from Boban – we'll capture that and we'll get an answer. I think that's relevant to this discussion.

BOBAN KRSIC: Anything else?

STEVE CONTE: So we're at the 15-minute mark until break, but then we have an additional 30 minutes with potentially Sam afterwards. So really, if there are no more questions, I don't want to waste Sam's time. If there are further questions when Sam's not here, we can always capture

them, but I don't want to force this to a close either. If you guys have more thoughts and questions.

So just for the record, we're having off-mic discussions about publicly discussable data versus data that might fall into more sensitive subjects, and so we're talking it through. At this point, the room feels more comfortable keeping us off mic, so we're going to stay off-mic for a few minutes while we discuss this.

Okay. We're back on the record. There was some discussion off-mic about the concern that the questions that might arise and have been arising and within IT and within ICANN's digital security, even though the intention is to look at the way that ICANN manages – the SSR of ICANN as it relates to this management and impact of unique identifiers.

However, there's no clear line that we can see, the Review Team can see, between some of the infrastructure that is running ICANN as a business and some of the systems and structure that's running the parts of ICANN that helps to coordinate and manage the unique identifiers that is in its mandate.

The concern was that if more pointed questions were asked, that it could be viewed or construed as going down inappropriate paths, and so the discussion was had that – and the request was made that if there was a way that ICANN can provide some kind of list or guidance on what systems or what processes – and please feel free to correct this wording – could help guide the Review Team to be looking in the right areas and asking more succinct or distinct questions surrounding the SSR portion

of the ICANN infrastructure. Is that accurate to the Review Team? Did I summarize that correctly?

JAMES GANNON: Yes, I think that falls in line with the off-mic discussion that we had. Essentially, we're kind of feeling around in the dark about what systems we need to be paying attention to and what processes we need to be aware of. And without wanting to go outside of our remit, we would like ICANN staff to be able to bring us something that is proactive. Because you guys know your network and you guys know where everything sits, and what is Internet-facing in terms of the DNS-facing, and what is traditional back office back office. And anything that you can do to present to us on what that is and what processes and procedures that you have that support that with regards to maintaining security and stability of those systems, it's exactly what we need to progress in that area.

STEVE CONTE: Can anyone else from the Review Team also chip in and say that you agree with this or disagree with it? I'd like to get some sense of the team.

NORM RITCHIE: Yes. I guess, yes.

STEVE CONTE: Alright, we're going to go ahead and pause the recording, please. We're having a fire drill. It's 2:00. It was supposed to be 3:00. Surprise! We're going to stay here until our fire warden comes.

All right, I apologize for the delay. We had a fire drill. The good news is everyone is okay, and the fire drill was successful. The bad news is we lost some time here. So we were actually just recapping a conversation about how ICANN can help inform the dialog to the subteam on looking in the right place and asking the right questions that doesn't feel like it's impeding out of the remit of this Review Team.

I was going around, and Boban I think was making some comments, or somebody was. James had already supported it and commented. We'll capture all that, but I think Norm and Boban wants –

NORM RITCHIE: Yes. I agree.

BOBAN KRSIC: I agreed also.

ZARKO KECIC: Me too.

BOBAN KRSIC: Okay.

DENISE MICHEL: I agree. Does anyone not agree?

UNIDENTIFIED MALE: [inaudible]

STEVE CONTE: We have consensus. We don't have any disagreement, so Sam, we'll capture this question and work with you and your team to provide a high level first entry into this conversation. As mentioned, it'll probably take some iteration, vetting internally, and so we'll get this back to the subteam as soon as we can, but to make sure that we can provide you guys with good information that'll help support your dialog.

JAMES GANNON: If I can potentially put out an ask that if we could get that before Abu Dhabi, I think that'd be fantastic.

STEVE CONTE: So noted, and I can try. It's not very long, and there are dependencies on people that I can't guarantee their availability. But I will make sure that we try to work towards that timeline.

DENISE MICHEL: Staff generally have a lot to do, don't they, before an ICANN meeting already. Yes.

STEVE CONTE: Any other questions for Sam at this point then? Otherwise, we're going to give him back an hour or a half hour.

UNIDENTIFIED MALE: [inaudible]

STEVE CONTE: Yes. Sam, thank you for joining us. We'll be in contact and we'll make sure that we get the right information back to the subteam.

[inaudible] I apologize for that. We are at break now, but I didn't want to leave Sam waiting on us because of the fire drill, and we were mid-conversation on that.

So our next scheduled person, team would be at 3:30 back with Jim Caulfield and Xavier. So you actually have between now which is 2:35 and 3:30, about an hour to either take along break or continue you guys' discussions earlier on the Google doc that you've been working on. I'll leave it to the team to figure out what to do.

BOBAN KRSIC: What do you think about if we have a break until 3:00, and then follow up with [us] catching the outcome of the last topic? That was [inaudible] compliance, Registrar Agreement or the compliance topic, and security incident management for half an hour, and then move forward with the last one, risk management, business continuity management, and then finished with AoBs and closing remarks. James, do you want to leave us 20 minutes?

JAMES GANNON: I'm just looking at Uber now. I need to leave in 25 minutes, so I probably will not be here for the next session. So thank you, everybody. I thought it was very productive, and I think we got a lot done and we have a lot of follow-up to do. I'm very interested to get that report and the drafting done on what we did and our potential recommendations. And I will not see you all in Abu Dhabi, and I will see you all in Brussels next year, unfortunately.

STEVE CONTE: From a logistics perspective, would it make sense – and I'll just leave this to the team to decide – to delay the break and maximize time while James is here? And then once he has to go, start the break at that point? Okay, so is the next part going to be more work on the document then?

So we're going to go ahead and pause the recording then because it was extremely auditorily exciting to watch you guys type, and I don't think we need to have the recording on, unless there's any opposition to that. We can insert in post the sound of keyboards typing for that period of time. And then when James leaves, we'll take a break. Please mark it when we do take a break, Yvette, and then we'll come back and have on-record conversations.

UNIDENTIFIED MALE: [inaudible]

STEVE CONTE: Thank you.

NEGAR FARZINNIA: Boban, if you're okay with it, we'll go ahead and start the recording now.

BOBAN KRSIC: [Yes.]

NEGAR FARZINNIA: Thank you. All right, everyone. Thank you so much. We are back in session now. We have Xavier and James here with us for our last session of the day to go over business continuity discussion. And with that, Boban, I'll turn it over to you.

BOBAN KRSIC: Yes, thank you. Welcome back. Same people, same place, other day, but same topics. And the idea for today is we had some business continuity-related questions that we would like to address, and also some regarding risk management, but there are only a few because we talked yesterday about the risk management processes and procedures here at ICANN, and we got a good overview of this topic.

Business continuity. We talked yesterday about the process itself as a part of the risk management, and we talked about business continuity plans and business impact analysis, and also of the topic of exercises.

And a question that we have is not the frequency, but how you provide these exercises.

Are these table rounds? It's only a theoretical part, or do you really execute it as a scenario, you define a scenario and then go through step by step active through this scenario to ensure that you have really everything in place that they can be 99.9% security would also work if you have an incident or a disruption for a long time. So how about these topics of exercising?

JAMES CAULFIELD: The question is, do we test the programs that we have? And the answer is yes. Today saw a test of our emergency response. We had a fire drill. That is part of an event, and one of those is to evacuate a building. And so I was almost wondering if that was done on purpose, but if it wasn't – I hadn't heard anything, so I would say no, honestly. But I think that that was a good indication that the first phase of resiliency management.

The second step is crisis management, and we do practice our crisis management. We've had a couple of tabletop exercises, and we've actually had a couple of live exercises. So we know that the crisis management team process works. And as far as [disaster] recovery goes, we do failover our systems to make sure that they're working properly. Most of our systems are [active at] dual datacenters, so you would expect that we would have a high confidence being able to restore our datacenters, and that is the case. We do test that annually.

UNIDENTIFIED MALE: [inaudible]

NEGAR FARZINNIA: Mic please.

ERIC OSTERWEIL: I'm sorry, I missed the first part of that. What do you specifically test annually? [inaudible]

JAMES CAULFIELD: Disaster recovery for IT systems.

ERIC OSTERWEIL: Okay. Thank you.

BOBAN KRSIC: Do we have another point to address? Especially in this context.

ERIC OSTERWEIL: Yes. Testing the DR annually, can you go into any level of detail on how you test DR? And do you do anything besides DR on an annual – like tabletops that includes attacks, outages, they include critical infrastructure like power outages and stuff? I'm just kind of throwing random things out. I'm just wondering if there's a little more detail.

JAMES CAULFIELD: Sure. We're not so much interested in the cause as the effect, so whether you lose your datacenter due to a power outage, terrorism, fire, we don't really care as much as what the result is, is that the datacenter doesn't work. And so we shut down one system and see if it works seamlessly with the other active center. If it continues working, then [it gives us a high confidence] level that if we were to lose that datacenter for any reason, the other one would keep going. And I think it's fair to say – there is one here in California and one in Washington DC. I don't think that's a secret. I think it's actually public information, so fairly public.

So we have two datacenters in different parts of the country, which I think is an important thing. And having worked at organizations who have one across the river from the other, that doesn't work. But I think here, it's clearly on the other side of the country, not subject to the same electric grid, not subject to the same kinds of natural disasters. And they do [fail out] the system and see if it keeps running on the other side. And it does.

ERIC OSTERWEIL: Thanks. That's great level of detail. And you're absolutely right, I don't want to get down in the weeds about anything more than we need to. But do you guys do any sort of tabletop exercises that sort of – in the event that you have a full disaster and the datacenter is inoperable, some of the processes are a little more clear cut of what to do than a ramp-up in an attack, an availability problem. Do you tabletop sort of triage emerging threats or emerging attacks or something?

JAMES CAULFIELD: Not that I'm aware of, but I've been here since April, so we might have to check on that. Xavier, do you know? No, I don't know, and I do know that there was some discussion a round that. Being the head of risk management, I was informed that they were planning to do something [along those lines.] But they could have been done within the last year, I just want here and I don't know about it. I'd have to check.

XAVIER CALVEZ: Not changing anything to the response that James provided, what we are trying to set up are exercises that combine several different events in the same timeframe. So you have a cybersecurity issue. Fine, how do we handle that? We have a plan for that, we have a physical security set of plans, but then what happens when you have both at the same time or three at the same time? It's the perfect storm issue, and so that's something we have n to yet done, but this is in the next steps of the tabletops that I know we want to try to simulate, is combining different events into one timeframe so that we can actually test the resiliency of the processes together.

So reacting to cybersecurity without having access to the building because we had to leave, therefore the computers are here. How we do that is then the next level of complexity.

JAMES CAULFIELD: We've done that for the CMT, as you're aware, but in that case, we didn't actually fail out the system. That was just tabletop.

XAVIER CALVEZ: The tabletop exercise that we did for the crisis management team, we had a number of scenarios and it started with an employee issue, and then it translated into social media issue about the employee, then it became a physical access to the offices because of the employee issue and potential threat as a result of social media.

So we tried to combine the events and sources of risk and effects of risks, test our communication and recovery plans across those. But it simply illustrated that we need to do more of those, and maybe more comprehensive, and sometimes also in situ type of exercises.

So our physical Security Team and Cybersecurity Team are looking at – without trying to be too specific – if we're all at an ICANN meeting – which you all know what it looks like – how do we handle a cybersecurity event and an office security event?

If you take an example, that example and push it a bit out, at an ICANN meeting, all the officers are at the ICANN meeting. The officers are those who can approve payments, so if you have everyone in one location in one time zone that's far away from the U.S. for example where there is an office issue that needs to be addressed, how do we handle that?

That's the type of combination of events. So the worst case scenario is what we're trying to formulate to test on the tabletop exercise, the processes.

ERIC OSTERWEIL: That was great. Thank you very much.

BOBAN KRSIC: I have an addition to this question. Your business continuity plans. What was the scenario that you used for the concrete plan? Do we have it in place? Only for systems that are related to SSR DNS.

JAMES CAULFIELD: Sorry, again, it's the results. We just pull the plug on the system. I don't care about the scenario. Whether it's kerosene fire, earthquake, all we know is the system doesn't work. So if the datacenter here goes cold, do they continue working outside of Washington DC? The answer is yes.

You can spend a lot of time thinking of a lot of different scenarios, but the end result is the datacenter goes down. Does it automatically fail over to – doesn't even fail over, it continues operating because we're active active on most systems. We do have some hot standby. So you pull the plug, there's a hot standby. [Do they] pick up in DC? Yes, they do. Pull the plug in DC. Do the hot standbys in LA pick up? Yes, they do. But the most important systems are active active. Which is great, because it's a high-level resiliency. You wouldn't expect to have a problem. You could always have a problem. There could be [inaudible] you didn't realize, and so that's why you test. You just literally turn off the system here. Does it work over there? Yes. [inaudible]

ERIC OSTERWEIL: Just to sort of follow through, because I think these were kind of related questions. I think the question is when there's something happening

and you're not sure what it is, you don't know whether to go to a DR mode or not.

For example, sometimes in a tabletop, you'll say, "Okay, suppose you'll now get the call that says this, and maybe have a database and you're worried about it going split brained. You don't want to go split brained, you want to fail over first."

So I think part of the question is when you build these scenarios, you're exercising your operational processes to know, will you do proactive measures? So it's at a higher level. I think that's kind of one of the –

JAMES CAULFIELD: Okay, that's a fair point. We do take into account in our processes decision making if we're making declensions on those kinds of things. So there is a discussion around that. In that sense, it's [sensitive] who would make decisions and how does it get declared that it's an emergency, and that kind of thing. We discuss it – for example, the crisis management team would be making those decisions with the IT team management.

BOBAN KRSIC: Regarding the crisis management, you're talking about teams. The roles and responsibilities, are they documented, known to people? What are [their role and] activity, which to have in their role?

JAMES CAULFIELD: Everyone's roles and activities are – there are plans for emergency response, for crisis management, disaster recovery. We have documents on that.

ZARKO KECIC: Do you have priority services that you're aware of, or all services are the same?

JAMES CAULFIELD: No. They are prioritized, and the most important things are – if we're talking systems – are active active, and then those are the more important things. And less important things might be hot standby. But that's things I don't think – I'd have to check with the technical side, but I think all the DNS-related things are active active.

XAVIER CALVEZ: I don't know if you've had conversations with LV McCoy for IANA, but the – and we can elaborate further on that point. Which are the services that are prioritized? I know that what we call DNS off switch is the L-root server systems, the IANA functions systems that are specific to IANA functions are part of that.

As part of the business impact analysis that we've done a couple of years ago, there was also those services that are considered most important for communications, and some that are part of that prioritization. So we'll circle back with the IT Team that mainly manages that, but also with the IANA Team to verify what is the level of priority [inaudible] but I think that's a very relevant question.

JAMES CAULFIELD: I would just reiterate, those are going to be active active, so I don't know how much more priority you can put on it. They've got two places running at the same time. That is the highest level of readiness.

ERIC OSTERWEIL: I was about to say this, I'm not sure if I should address this to you all or I missed talking to one of the earlier groups, so forgive me if this is sort of [inaudible] but one of the reasons that the active active – it's very great, so not to diminish that at all, but one reason why it gets more nuanced is if you have like an insider threat problem and you do need to offline the site for some particular reason, so the physical infrastructure, the cyber infrastructure is still working, but there's a reason that you follow business process to say, "The site is going offline now" or work on transfer operations on something like that. So this is the question, I'm not sure if it's appropriate for you all, but do you have an insider threat program? Training, briefings, processes, something like that?

JAMES CAULFIELD: It is part of our – especially on the security operations side.

XAVIER CALVEZ: As part of that exercise, one thing that's fairly straightforward that's logical but that contributes to it is the access rights that are – the IT team has been putting a lot of efforts over the last few months on reviewing in detail the access rights of every staff member against job descriptions and what are the needs and so on, so that we have a

"clean" database of access rights to systems, and you only have access to what you need to and no more. As well of course of the processes to update those access rights when positions change or status of people changes. So that's one of the things that is fairly standard process, but that's part of those exercises.

ZARCO KECIC: You mentioned a couple questions ago how you plan communication with professional organizations and media and authorities in case of disaster or event.

JAMES CAULFIELD: Sure. That's certainly part of the plan. Communications is documented in our plans: bulletin, EMT, disaster recovery. You may be aware we have a large professional communications team here at ICANN that would be also involved in that kind of thing. So it's clearly recognized as part of the plan.

ZARCO KECIC: But it's very tailored to the type of event, so you will of course communicate through the employees if it's an office-related function, including possibly contractors. When we did the tabletop exercise, it was very clear that in the decision making process of what actions need to be taken as a result of an event, including communications along with HR and legal into the small group of decision makers to define the actions resulting from an event, was important because that aspect of

communication is to be adjusted and tailored to the event and to potential future impacts.

So the Comms Team is involved at the core of the decision making as to what action needs to be taken as a result of an event, including communications either internally or externally or both on a given event. Of course, not everything receives the same treatment from a communications [standpoint].

JAMES CAULFIELD: That's clearly stated in our [business continuity].

DENISE MICHEL: Following up on that to give us a bit more context, would an incident such as the DDoS attack that took out several root servers, would that be the type of event that would trigger this? In other words, is this only relevant if it's affects direct ICANN operations or L-root, or does it also apply more broadly to the broad root zone system and more broad DNS emergencies or problems? Does that make sense?

XAVIER CALVEZ: I think it makes complete sense, and I will speak simply from memory of the more recent events that occurred. There has been fairly regular, and I think it's either every six months or every year, tabletop exercises done across the 12 root operators on resilience to simulated attacks. I don't know how broadly it has been communicated, but I know it has been happening and I know it was public because I've seen public communication about it.

When there has been an event, I know that we have at the minimum communicated with the amount of details or lack thereof that could be communicated publicly how L was impacted and our knowledge of potential impact across more than just L. But then we will need to, as we have, coordinate communication with each of the operators because at least as of today we are not necessarily the voice of each of the 12 root operators as to the status and resilience of their own root for a given event.

So I think everyone communicates about their own status to the extent that they can and desire to. What I know we have tried to do is at least check among all the root server operators what their status is and so on. What gets communicated out of that exercise of checking how everyone is doing I don't have that in mind. We can check with the DNS Ops Team and Comms as well if there's a standard protocol of communication on that, but I don't have it on top of my head.

BOBAN KRSIC: Any other business continuity related questions? No? Then I would like to come back to risk management. We heard yesterday a lot of your process and your risk treatment strategy. My perspective on this is that we are talking only about one strategy and that [involves] mitigating with controls. Now do you have also other strategies in place like acceptance of risks or something else, or do you [inaudible] strategies only we will be sure to have the right control in place to mitigate the risk and the possibility of risk?

XAVIER CALVEZ: Some of those risks could be [inaudible]. That's something that, by the way, resulted from our review with the Risk Committee of the Board. And we need to be careful that some risks may be different, depending upon whether the transition happens or not.

Then who is the risk owner? The risk [delegate simply usually so] I'm the owner of some of the risk and then someone in my team is going to be the delegate because they're the ones maybe carrying out the mitigation [plan]. So risk owner and the delegate of the risk owner [we'll assign].

This is the year during which the risk was identified. You can see that the rectangle within a red outline is what I did to hide information. I didn't hide all this, for example, which is illustrating how we rate the risk. This is when the risk was identified and first made it into this spreadsheet.

Then we assess the likelihood from 1 to 5, 1 being low likelihood of risk and 5 being high. Impact or severity of the effect of the risk occurring from 1 to 5. Then inherent risk: we basically multiply the likelihood by the impact. So highly likely, high impact, very high inherent risk.

Then we try to rate the effectiveness of the controls or mitigation that exist that are in place. We changed that rating last year to make it from 1 to 3 to 1 to 5 because we needed a bit more depth in that rating. So now it's a 1 to 5 rating. Then to get the residual risk, we simply divide the inherent risk by the rating of the control effectiveness. That then results in the residual risk.

So high likelihood will result in a high inherent risk, high severity as well. Low control effectiveness will result in a residual risk that's high as well. And of course, likelihood and severity, if it increases, it increases the residual risk. Control effectiveness, if it increases, it reduces the residual risk. And we track both before and after.

UNIDENTIFIED MALE: Can you explain what is control effectiveness?

XAVIER CALVEZ: It will vary based on each of the risks, but it's basically what do we think are the mechanisms or controls that are in place in the organization as a whole or specifically in the function that manages the specific risk to mitigate the impacts of the risks. If I try to take an example, and again it will be very dependent upon the risks, in case of a disaster, for example, the control effectiveness is (from memory) I think it is considered medium because we have zero control about the occurrence of the event like an earthquake but we have controls on the mitigating actions that we have in place. So for example, we have two data centers to mitigate that.

Now, I will admit there's no question about it, this is very subjective. Whether it's 2 or 3 [is a bit of debate on this], to be honest. But as we are making progress and educating more and more the organization about measuring this, we're starting to be a bit better at this. We started developing criteria to say what is a 1, what is a 2, what is a 3. I think that we'll need to reassess how effective that rating is. It's not an exact science. As you know, it's not meant to be. You can rate

everything you want. You can put a lot of arithmetic behind the risk assessment. It still remains subjective. All of these ratings are subjective. There's not a very binary set of criteria to decide what it should be. It's a guess, and it's also something we monitor over time for changes.

UNIDENTIFIED MALE: I would just point out that our definitions of what the controls are helps you have some context. It isn't just, "Hey, what do you think it is, a 5 or a 1?" It's based on is it documented, is it repeatable, is it observable? There are some criteria around to help guide whether it's a strong or weak control or [mitigate].

XAVIER CALVEZ: Right. We've tried to introduce those criteria to help the risk owners do a better job at evaluating control effectiveness because we could see that it was very subjective and dependent upon the expertise or experience of the managers.

UNIDENTIFIED MALE: Right, so it remains subjective, but there are guides to try to calibrate across the control effectiveness.

XAVIER CALVEZ: : So if I continue down the line, then that results into a ranking. Ranking in the sense of what's the highest residual risk to lowest. So we rank that. We do two rankings. We try to keep a list of top risks. You may see somewhere – in the next document, you'll see it think that there's a top

13. Why 13? It's because we had ties in the ranking of the residual risks, so we took all the risks that were down to a certain rank and that was 13 risks rather than 10. But otherwise, we tried to stick to what are the top 10 risks and monitor those risks throughout the period even if their rating brings them below the top 10 because we want to keep visibility on those for a little while.

Since we do a quarterly assessment, the top 10 could change every quarter and sometimes they do. But we still wanted, because we had established that those were important risks when we do the update on an annual basis of the risk [register], we are trying to keep a certain constant look at the same top ten. Though, of course, if one risk comes up into the list of top 10 that was not before, that's exactly what we want to know. Suddenly there's a risk that has increased, and we should bring that up. We should look at it and define what we do about it.

So we follow the ranking. We had a slight change in methodology when I told you that we want from a 1 to 3 scale for control effectiveness to 1 to 5. So that's why we track the pre-March 2016 ranking. This is logistics and not very important anymore.

Risk decision is what I was saying earlier. It's the mitigate, ignore. That's where the risk decision is. [Remediation] plan is an overview of the remediation plan that the risk owner is in the process of either defining or implementing or both. Of course, that's something that we also provide visibility to the risk committee on.

We had as well a post-USG transition remediation plan should the transition occur, so that's a column that we're leaving blank now but that we felt was important at the time [inaudible].

Let me show you the second picture, which is simply the columns on the right.

BOBAN KRSIC: Xavier, before we go to the second one, thank you for this great [overview] and it helps really to understand your approach in risk assessment.

Two questions. Question number one: it's a qualitative approach, yes? So do you quantify it, your top 13 risks, and make some [inaudible] financial perspective? Because we know there are risks. Yes, we have to handle. As an organization, we decide to make reserves on the [inaudible] risk related to our, I don't know, budget or what else. That's the first one.

The second one: the tracking of control effectiveness. If you have controls in place that you plan controls to mitigate the risks, how do you track them? Are they implemented or not implemented, and what decisions are you making after that?

XAVIER CALVEZ: We do not create financial reserves specifically for a set of risks. You may know that ICANN has a reserve fund, which is the financial mechanism to ensure sustainability and continuity of the operations for ICANN. Which, by the way, you guys may want to be interested in that.

This is one of the pillars of how ICANN ensures the financial sustainability of its operations. They have the reserve fund.

There's a completely separate exercise. You have the primary of the information. I'm about to send for public comment a document on the reserve fund and the rationale for the reserve fund and what its target level should be. It currently is of 12 months of operating expenses. But in that rationale, there is an exercise of risk evaluation that entered into it. So looking at it from a risk perspective, the answer to your question is the reserve fund is what we have in place for risks.

The control effectiveness and the mitigation plans, the more we have been ramping up in perfecting or improving this tracking mechanism, the more we have narrowed down the weaknesses of our monitoring on the amount of understanding, visibility, and effective control we have over the mitigation plan and measuring the progress, measuring the status, and then measuring the effectiveness of those mitigation plans.

We focused a lot of efforts over the past couple of years on identification and monitoring, which helps us understand that now the next stream of efforts is on adequacy of the mitigation plans to the risk identified and then effectiveness of the mitigation plans.

I won't try to hide the fact that for us this is an area of improvement. As we have [inaudible] risk owners, the quality of the information that we're collecting through this process is also as varied as the experience of the risk owners in adequately identify the plans, documenting them, and making progress on them. So this is something we need to make more efforts on. We have discussed that we the Board and the Board

Risk Committee, and they're definitely sensitive to that because we've made them sensitive. We are sensitive to it, and we want to make efforts [toward that].

UNIDENTIFIED MALE: Thank you.

XAVIER CALVEZ: I'm trying to share now the rest of the document which is simply the rightmost columns of the same document. Then I have the third document quickly to show you to finish on that. Great. One was working; the other one doesn't. Same exact document. Same picture. Let me try again. I'm not sure why it doesn't. Okay, sorry. It doesn't work. I have the picture here. Oh, is it me doing that? Right, but I was trying to share the other picture, but it's not working. I've tried a couple times. It's simply not working.

UNIDENTIFIED MALE: [inaudible] I've got to load it. Stop sharing

XAVIER CALVEZ: Stop sharing? Okay.

UNIDENTIFIED MALE: [inaudible]

XAVIER CALVEZ: Okay, looks better. [Nope].

UNIDENTIFIED MALE: There's a beautiful picture in Adobe.

XAVIER CALVEZ: Yeah. So the columns that we were not seeing, I'm going to go quickly over them, actually pertain to the mitigation plan. We were looking at risk decision. Then there is the remediation plan, whether it had a USG transition impact, the mitigation status: Is it on track? Is it late? Is it [inaudible]? The key success factor for the plan to be effective, then the KPI. So basically, what success looks like once we have implemented the mitigation plan. The KPI to measure progress toward the implementation of the mitigation plan. And the last column is reason for change, if there has been a change in the mitigation plan. So it's just trying to document the mitigation plan. So that's those.

Let me then try to show you the slide that we provide to the risk committee of the Board every quarter. Share document. Hopefully, this one is going to work. I don't know what I'm doing here.

We provide PowerPoint materials to the Board members prior to each committee meeting. This is a fairly standard slide that we provide which simply gives as you'll see – I'll switch microphones. You remember the columns of inherent risk which resulted from likelihood and impact multiplied by 5.

So we have here, this is more the top risks. So this is the top 13 from memory. We've aggregated together, averaged the inherent risk rating

of the top 10. The maximum is 25 of course, 5 x 5. So before we do anything about these risks, this is what they look like in likelihood or impact. Then we average out the control effectiveness for each. And then we average out the residual risk for each. By the way, if you divide this by that, you don't always get this because this is an average.

Therefore, we show to the Board what changed about the nature of the risks, what changed possibly about the controls that we've put in place, and what changed as a result [of] the residual risk. You can see that from 1 to 4 there has been relatively little change to the control effectiveness, and two of the risks by memory increased a little bit.

UNIDENTIFIED MALE: One.

XAVIER CALVEZ: One. Here, we're explaining what changed about the risk. I cover that because then we're listing some of the risks [inaudible]. But that's an illustration of the reporting that happens on a quarterly basis. One of you discussed yesterday with Steve, "Can we just have a proof of life?" This what I was hoping could address the point. Thank you.

STEVE CONTE: I just want to remind you, Boban, the sensitivity of time and their availability. So we are at time for that, so I don't know what your availability is.

XAVIER CALVEZ: We can make ourselves available until 5:00 p.m.

BOBAN KRSIC: Any questions? Thank you very much also to show us what we can [inaudible]. It makes it much easier to understand it, and it's a great approach how to deal with your risk management approach at ICANN in all aspects: how you identify new risks, how you categorize them, how you try to mitigate them, the reporting. I think the Board makes also the right decision from them what they see and give you the right resources you need. So thank you very much. It helps to understand this whole topic.

From my side, I don't have any questions related to business continuity and risk management. Thanks again for your time here and wishing you all –

UNIDENTIFIED MALE: Thank you for your time as well.

STEVE CONTE: So we had a period of time where we had no audio. Is there anything that we should recap in the dialogue just to make sure we capture it in the audio and before they walk away that everyone is nodding and agreeing with what dialogue took place before we got the audio back up.

UNIDENTIFIED MALE: I don't know when [inaudible].

STEVE CONTE: It's sometime just before I walked in.

NEGAR FARZINNIA: I don't know exactly when we lost audio. Right before you had walked in, we actually were spending time looking for documents. I believe that's what Xavier was doing, so there wasn't a lot of conversation to be captured.

STEVE CONTE: Was there any specific…?

ERIC OSTERWEIL: No, I don't recall any asks. I think what we talked about was in regards to disaster recovery. We talked about failover backup sites, active active, that sort of stuff. I think it was mostly getting Q&A done around that stuff. I don't recall any asks.

STEVE CONTE: Okay, thank you. I just want to make sure because we are going to be using the audio to capture the questions, and I wanted to make sure we captured anything that might have been lost with no audio. Okay, thank you. Thanks, guys.

NEGAR FARZINNIA: All right, Boban, back to you. I [think] we have the Any Other Business section and Closing Remarks. So back to you guys.

BOBAN KRSIC: Yeah, thank you very much. So, team, what do we think about this session when you compare them to yesterday's forum. Was it better? Yes? Much more detailed? [inaudible]

UNIDENTIFIED MALE: Thanks for the feedback too. I think that's the kind of thing that really helps. It's challenging for all teams to expose themselves. That sounds awful, and I don't mean it in an awful way. But it's challenging, so having that first presentation and then feedback from that I think helped. I hope to continue that dialogue with the team so we can make sure that the discussions with ICANN staff remain fruitful in a constructive and positive manner. So thank you for that feedback, [inaudible].

BOBAN KRSIC: So to recap this session, I propose to write it down in a document so everyone can write down in a Google Doc some thoughts about the business continuity, disaster recovery planning and so on. I have some [remarks] about risk management. Then we should try to finalize a document to give them a structure and to identify the open issues to address them to the ICANN staff to get more details on any questions.

Any feedback from the team?

ERIC OSTERWEIL: I think, yeah, that level of transparency was extremely helpful in helping us feel comfortable about what's going on without having to get up in the grill of everything in detail. So I think that was really useful. I think there are a number of things that we may just need individuals on the team [iterate] in the Google Doc in order to get on the same page for the discussion that we should have. I think we should do that expeditiously I think. At this point, that Google Doc represents a pretty significant step forward as far as work product from the team, so I think the sooner we can put that into some shape to start talking about the better. Because I think this was a huge commitment of time for the ICANN staff for various people to take time out of their day and answer our questions and [iterate] with us. I think we want to show good faith and get it back out quickly.

[DENISE MICHEL]: That's called equalizing. Just use the situation effectively.

BOBAN KRSIC: Great. Any other business?

[DENISE MICHEL]: Over the course of today and yesterday, obviously we had so much and items that came out as a result of these conversations and discussions and questions that require follow up from the team. There were also some questions e-mailed out throughout the two days. I'm going to have my team work on putting together a list of all of these action

items, going through the recordings and listening in to everything that we think should be captured.

What I would like to do is [have us follow] our normal process, send it to the subteam, have you guys approve to ensure everything has been captured properly. Then once we have your approval on the list of questions and action items, we'll take it back to the internal team, distribute it amongst the SMEs to try to get some timeline from them on responses they will provide to you.

Just please do note that because it is so close to ICANN60 and with the workload that everybody has on their plate, the responses may be delayed until post-ICANN60. But we'll do our best to get everything to you as soon as we can.

BOBAN KRSIC: Thanks in advance. So then I would like to say thank you to all of them, to team members, to you Steve, Negar, Yvette. Thanks to all. It was great having you all here to work on them, two intensive days with a lot of stuff. I think we are a big step moving forward to our goal. Thank you from our side for this.

[DENISE MICHEL]: Thank you, Boban, for organizing this. I think it has been really fruitful. Thank you Yvette and Negar and Steve. Really appreciate your support on this.

UNIDENTIFIED MALE: Please pass our thanks [around to all] [inaudible].

UNIDENTIFIED FEMALE: Yeah, to the staff. Yeah.

[STEVE CONTE]: Will do. Absolutely.

NEGAR FARZINNIA: Absolutely. It has been great having you guys here. I think it has been a productive two days, and looking forward to our discussions in Abu Dhabi.

BOBAN KRSIC: With that, let's go ahead and stop the recording and officially end proceedings.

[END OF TRANSCRIPTION]