GEOFF HUSTON: And Brian. I see Brian listed here. Okay, thank you all. Let's make a start.

Kim, I've asked you to join us because there are a number of questions here that you can either answer or point us to the relevant documentation or give us any other observations about these. In some ways, you know, most of them are no-brainers, but it's kind of a formalism here that we just like to understand a little bit more about how you and your group look after the root zone of the DNS and the relevance of the explorations around the boundaries. So thank you for coming and with that, maybe we should dive right in.

I grouped up these questions thematically, Kim, so the first is kind of about label management and the next around mechanics of delegation, and the last set around DNSSEC and rolling keys. So if it's okay with you, perhaps if I throw it over to you with the first sort of set of questions about TLD label management, looking at where are the guidelines and constraints about the root zone of the DNS and how is that managed in terms of which groups, [without] the betting, etc.

KIM DAVIES: Thanks, Geoff. Firstly, thanks for the invitation. I think, as IANA staff, we really appreciate being involved in these discussions when we can to try and explain our operation procedures and hopefully it helps inform the outcome of this process.

That said, what guidelines and constraints govern the labels that are placed in the root zone of the DNS? For example, single letter character

domains, either an ASCII or the Internet equivalent to [inaudible], two-letter code other than those defined in ISO-3166 [inaudible].

So generally, we define what is allowed in the root zone by inclusion rather than exclusion. So single character domain names from an IANA perspective are not disallowed. However, at this time, there's no past eligibility for them under the existing policies.

So what are the policies? The policies are the relevant ccTLD and gTLD policies which to date have not been overlapping or conflicting, but that is something that we sort of rely upon, that the gTLD policy doesn't permit a ccTLD and vice versa.

For there to be a pathway and eligibility for single character domains, for example, that would need to come as a newly ratified policy through the ICANN policymaking process. So how that would play out is that the GNSO, ccNSO would come up with a policy that would get ratified through the relevant organizations, it would go to the ICANN Board for approval, and then staff would be instructed to implement that policy by virtue of a Board resolution.

That's the high level. Do you want me to drill down into that any further or is that the kind of level of detail you're looking at for this discussion?

GEOFF HUSTON: No, I think that's entirely appropriate. I had just one minor thing in my head, which sort of crosses over with Unicode. Under the existing policy framework, a single Unicode character would encode as a multi-character [LD string]. Would that be allowed under the existing

framework, because it's actually a single character in its Unicode form and a multi-character stream in its Unicode encoded form?

KIM DAVIES: Right. So, I mean, there's no perfect rule here, but generally, we try to create… [Consider] Unicode characters as first-class citizens. So if it was a single Unicode code point that happened to be encoded as multiple ASCII characters, we'd still consider it a single character domain.

I think in terms of what the constraints are, there's no pathway that any single Unicode character point could be a TLD today simply because under both the prevailing gTLD policy and the prevailing ccTLD policy both rule out single character domains whether it's the ASCII encoding or the Unicode representation.

So I think, you know, there's no prohibition on single character domains as I mentioned earlier. It's just that there's no pathway, there's no scope in any of these new policies to allow them either. And that's, and I think by accident, I think the community has expressed reservations about single character top-level domains that they've implemented and today is known as a prohibition, but simply through the policies stating that eligible TLDs must be a minimum of two or three characters, or whatever the case may be.

GEOFF HUSTON: Sorry. I'm reaching for my [inaudible]. Thank you very much for that, Kim.

In so answering, you've also answered the second question that you're pointing to ccNSO and GNSO policies. I had just one sort of side question about that, which is underlying this is the RFCs coming from the ITF, and in particular, the so-called LDH rule.

KIM DAVIES: Right.

GEOFF HUSTON: Are those constraints from the ITF over-arching constraints, or if the policy communities in ICANN want to go into different places, can they go there? Where's the boundary line, or haven't we explored that?

KIM DAVIES: I think this is an area that has lacked full exploration. I think in a practical way, what happens is that usually as policy is developed, ICANN staff including staff who perform the IANA functions are usually involved in some capacity, either as an invited expert or somehow convey the nature of the discussion.

So at that stage of policy discussions, if it was touching on these kinds of issues where we [can think] forward down towards implementation, we could try and flag them for consideration via the appropriate mechanisms. It's important for us not to influence policymaking, but I think it's fair to say that we can at least bring issues to their attention if they hadn't otherwise considered them.

Now, they may consider them and decide to proceed with [inaudible]. I think that's within policymaking community's remit and that's not something that is often constrained. If there was to be a policy that was implemented that we thought contradicted RFCs, I think the way policy is usually implemented is that, like I said, it was ratified by the ICANN Board, they would recommend it to staff for implementation.

At that point, we would flag implementation issues and in the scenario you proposed, I think that's rare. If it had gone that far, we would identify that here is a conflict between proposed policy and RFC or other technical standard, and we would provide a recommendation on what to do, and that could be that the Board refers it back to the community to be any number of things. I think it's highly speculative because this hasn't really played out that far today.

And then if the community ultimately said we needed to implement it this way, I mean, I think I'd explore territory. I think without any contradiction from other parties, I think we would proceed on the basis that it's now the new ICANN policy. Like we said, I think that's the limit of how far I can see into the future. I think it strikes me as relatively unlikely we would get that far down the path because there are a number of avenues throughout the process to raise the issue fairly early on and those that wish to interject into the discussion would have an opportunity to do so prior to that time.

GEOFF HUSTON: Thanks, Kim. I think we're almost [inaudible]. We're scooting along really well.

I wanted to actually ask about Unicode and the IDNA rules, and I note SSAC has got itself wound up over this. I think it's got [EU] in some other script language.

I suppose my particular question is do you think the IGF in its IDNA documents has given you enough context to work from? Or is Unicode a gray area in terms of the staff's interpretation of guidelines?

KIM DAVIES: It's hard to know where to take that question. I mean, in terms of the IDNA specification itself, I think it's relatively straightforward. I don't think I have any personal concerns or wearing an IANA hat, any implementation concerns as it comes down to the wire format and how things are converted back and forth.

I think, to digress for a moment, there is an issue that one of the co-points was implemented in Unicode Version 7. There's differing opinions about whether it should be valid or not. The IAB gave instruction to IANA to restrain from publishing IDNA tables from Version 7 onwards as sort of a stop gap and that has now become sort of the permanent state of affairs because that work hasn't advanced in several years, so we're sort of stuck there.

I think from an IANA perspective, that's not too much of a concern, but that is something that one of the protocol [premise] side of the fence, an oddity that we have to deal with.

So that digression aside, in terms of implementing TLDs, I don't see that there is any problem with the specifications as written. But perhaps I

could ask the question back. Was there a certain scenario you had in mind in that question? I guess it's worth distinguishing that by the time it comes to IANA, the labels themselves [inaudible] well and truly decided. IANA isn't in the business of deciding what the labels are to represent a country or a territory, or in the case of gTLDs, anything else. They are vetted well in advance of them coming to IANA to implementation and then root, so any qualms about what label represents what, should have been ironed out prior to us seeing it.

GEOFF HUSTON: Some of this came from my own thinking, Kim, about the issue over the incorporation of emoji characters in second and lower levels in the DNS. And this larger issue that it is possible in the world of Unicode to create arbitrarily strange sets of symbols they encode into Unicode into perfectly conformant LDH. But in some ways, the Unicode itself is displayed in different ways on different systems incorporates display points that appear to be punctuation, etc., and I suppose my question really was are you effectively following simply what the RFC is saying which is not that descriptive in this area, or do you simply accept what comes through from the ccNSO and GNSO? Or what sort of criteria are you applying once it comes to you, if any?

KIM DAVIES: So today, the only way that IDMs are deployed in the root zone is by one of two avenues. One is that it comes through what is called the IDN Fast Track, which probably is a misnomer at this point. But essentially, that's the ccTLD eligibility pass. There, they have to demonstrate that

the code points are a meaningful representation of a country's name. So it's a fairly constrained set of requirements and there's consultation with linguists involved in that process and so forth.

In the gTLD application round, obviously, that is a [known] set today. There are 116, I think, that were applied for. I can't recall how many made it through, but they similarly had some eligibility process where they had to demonstrate the meaning of the term. There were reviews by experts as part of the review process in terms of evaluation panels.

So, in short, yes. By the time IANA gets them today, they would have gone through some vetting process either through the gTLD or ccTLD process, and we would assume that there is no conflict there. I mean, obviously, it needs to work in our system. Any eligibility per the standard would need to be met for it to even be inputted into our management system, so if it was nonconformant IDNA per the RFC, that would be a showstopper, but I can't conceive that if we get that far, with that being a problem.

So I think things like emojis are [nonstandard] because emojis are plain-as-day illegal according to the IDNA standard.

One development I will flag that might be of interest in the future is that there is ongoing effort in ICANN for many years to implement a so-called variance in the root zone. One aspect of that is this thing called the Root Zone LGR, Label Generation Rule Set. And what that is, is a conglomeration of rules from different languages that have been developed within those language communities. And what we perceive that will likely happen in the future is that that Root LGR will serve as an

additional set of rules that govern labels in the root zone. That's not the case today. It's still under development. But the Root LGR will likely, at some point in the future, be an additional constraint that might address some of the areas of concern that you have.

The Root LGRs are designed to be a somewhat conservative set of valid code points that further constrain beyond simply what is allowed by the IDNA standard, and also provides rules for generating variance so that, for example, we could conceive of a future where if you go to a Chinese language TLD for example, you might automatically get both the simplified and traditional version of that delegated together as some kind of bundle so that they would be treated as some kind of [atomic] set that is handled together even though they are multiple labels inside the zone file itself.

So that's not the case today, but that is what that line of work is foreseeing at some point in the future.

GEOFF HUSTON: Thanks, Kim. That's really helpful.

The next are three sets of questions, and I suspect that you have answered them already. But what it's both pointing to is that there is some degree of coordination in the two-letter country code with the work of ISO-3166 in the more amorphous space of generic names with the IGF and its special use names registry, and you have already covered the Unicode consortium.

So in both ISO-3166 and the IGF case, is there any particular input from staff or is this just a case of policy-based coordination?

KIM DAVIES: So in the case of the ISO-3166 standard, we have a relationship with them. It's probably, I've written some notes so I might just read through them and see if it covers most of what your concern was, or area of interest, I should say.

And I note that in your sub-point, you asked about exceptionally reserved, traditionally reserved, retirement, and so on. So I'll read to you what I have and it covers, I think, most of those topics. And then by all means, we can drill down into specifics as needed. I'll also flag that this is an area of active discussion in the ccNSO right now. In fact, myself and [inaudible] gave detailed presentations on this topic in Johannesburg, and I'm happy to share the slides with you, for example, if that's useful.

So what do we use ISO-3166 for? Well, the primary thing is we use ISO-3166 Part 1 altitude codes for a) determining eligibility and the label, so ASCII ccTLDs, and secondly, for determining eligibility only for non-ASCII ccTLDs. So essentially what that means is the ISO standard tells us that Australia is a country and AU is its altitude code. So we use both of those elements in the case of an ASCII ccTLD.

In the case of a non-ASCII, we would use it to tell us that Australia is a country, but what represents Australia in that non-ASCII label is a separate process. It doesn't involve the ISO standard.

We assume automatic eligibility for any two-letter code in that standard to be an ASCII ccTLD. For a non-ASCII ccTLD to be eligible, they both must be in the standard and then have their chosen string past the IDN Fast Track process, and that's administered in other parts of ICANN outside of IANA.

In terms of exceptionally reserved codes, there is limited eligibility for exceptionally reserved codes. There is an ICANN Board resolution that was passed in 2000 that actually governs this. So essentially, this Board resolution defines when an exceptionally reserved code may be a ccTLD, and as a result of that resolution, we bucket exceptionally reserved codes into three categories. The first category is codes that were delegated prior to 2000. In this case, so are delegated prior to that Board resolution that set out the requirements, we essentially grandfather them.

There's codes that meet the 2000 requirements. In this case, we consider them eligible because they meet their prevailing policy. And then we have codes that don't meet the 2000 requirements. In this case, we consider them in phase-out, so essentially, they are not eligible under today's governing policy and they're not grandfathered, but they need to be phased out to moderation.

In terms of transitionally reserved codes, they're also considered in phase-out. There's no codes today that fit that category, or the codes are transitionally reserved have already been phased out in terms of being removed from the root zone.

But whether it's, whatever the form of ineligibility, the phase-out process is essentially that we expect ccTLD managers to take responsibility for arranging the appropriate local mechanism to do a retirement. There's no prescribed approach. There's no strict guardrails much like the remainder of ccTLD policies. There's fairly wide latitude with ccTLD operators to do it in a way that comports with what makes sense for their particular local community.

And we operate on the basis of encouraging them to do it as reasonably quickly as possible given their constraint. Codes that have been transitionally reserved can be reassigned by the ISO-3166 standard. It could be put to use for some other purpose. Even though there is a target today that ISO will not recycle codes within 50 years, it's really just a target and the ISO community has told us that they may choose to reuse codes in a shorter period than that should they deem it appropriate. For example, there's actually not a lot of codes that are available for use, so if new countries are created and there's no good alternative, they may use a transitionally reserved code for that purpose.

There's no provision for delegation of indeterminately reserved code, so there's no mechanism whatsoever today to delegate those.

Let's see. And then finally, I'll turn back to our relationship with the maintenance agency. So I saw an approach to ICANN probably about ten years ago now to be a participant in the maintenance agency that administers the standards. They do this because we're a high profile user of the standard, arguably the highest profile user of the standard,

and therefore, from their perspective, having high profile users of the standard involved in its administration makes good sense for them.

The way we've [accepted] the invitation, but the way we've implemented it is that we contract a non-staff member to be our delegate to the maintenance agency. So that's somewhat at arm's length. We have a small internal coordination group of various people that are interested in ISO's administration. We're getting staff and that coordination group liaises with our delegate if matters come to the MA that require discussion. And also, I think most importantly, we've made the organizational decision to never vote on matters pertaining to adding and removing codes so that there's no perceived or actual conflict of interest that ICANN is picking winners and losers in terms of who gets a code.

So that's a rough overview of the notes I had under this question. Does that touch on all the areas of interest, or can I elaborate further?

GEOFF HUSTON: Just one final point, Kim. Do you and ICANN staff interpret this arrangement as effectively putting all ASCII two-letter country codes under the terms of this arrangement or just those ones listed?

KIM DAVIES: I would say from an implementation perspective, I think implicitly, this puts all two-letter codes under the remit of ISO-3166 adherence. And the reason I say that is that the ISO standard is dynamic. It does change.

I think once I did the math, one time per year a code is added or removed over time. That's roughly the average.

So assuming it's dynamic and assuming we're mandated to adhere to the ISO standard per ISC-1591, etc., if we were to only take a static view of that standard and say everything not in the standard today is out of scope, then whenever that standard changed, we would essentially no longer be adhering to it. We'd be adhering to an old version of it, and over time, we would drift apart and we could no longer adhere to it because one of these codes would have been assigned to some other purpose.

So I think as a practical implementation perspective, we consider all two-letter alpha code as in scope. I will note that the standard itself does [bring sense] a number of codes for user, defined users. So there's a hypothetical there that the user-defined purposes could be used to other purposes in a root zone context without fear they would ever be assigned as country codes. But that's also something that's being explored by the community to my knowledge, and it's not something that we have any policy to work with them.

GEOFF HUSTON: Okay. Thanks, Kim. Before I move on to the practical issues of NS and DS record management, does anyone else have questions to Kim about the mechanics of TLD label management for the root zone? No. Okay, let's move on some more. I was going to say "more prosaic" but it's certainly more down to earth issues, Kim, about the changes in what is the other

parts of the content of the root zone, notingly the NS records, the DS records, and even the glue records.

And I suppose the first question is procedural. When you get a change request – please delegate, please alter, please do this – and assumedly, you get some kind of assertion, "I am the authorized individual to pass you this," how can you tell the lies from the genuine requests? What practices do you use to sort of ensure you're not being duped?

KIM DAVIES: So, very good question. So firstly, we don't trust that at all. We don't require, actually, an assertion that the party submitting the change request has any standing whatsoever. In fact, it's not that uncommon that it might be, for example, a vendor of a registry or a consultant that they've engaged that is actually doing the submission on behalf of the formal operating institute that we have on file.

So the way we authorize changes is actually two primary mechanisms. One is positive consent from multiple designated contacts that represent the TLD managers. So regardless of who the party is that submits the change request, we today would look to the administrative and technical contacts both to explicitly authorize the change request. So in that instance, it's an automated procedure. Our root zone management system will e-mail those contacts that there is a request pending for them to authorize and they can log in to our web-based interface to review the nature of their request and submit an approval or a rejection of a request, and we require both of those to do that today.

The second thing when it comes to NS/DS records is that we cross-match what they've proposed with data that's in the child zone. In both those cases, an NS record change should be reflected at the apex of the child zone prior to submission for the root zone. And similarly, DS records should reflect DNS key records that are the child zone.

So by cross-matching those, as we evaluate a change request, essentially what we're doing is demonstrating that, firstly, they're accurate, that no errors were introduced in their submission, and also, that the request essentially has custody of the zone file itself because they've been able to make the requisite changes in the child prior to requesting the changes in the root zone.

So that's the primary mechanisms that we have today. I would say also, informally, today it's not a step that is fundamental to the process but nonetheless is there, that all root zone changes today are still manually reviewed, I think both by ICANN/PTI, and also by Verisign. So informally they're the sniff test there. If something looks unusual or weird, we would seek to clarify with the requester to see if there is something that we should know or understand the circumstances, why this anomalous-looking change is being proposed.

In particular, if there is a fundamental change like completely replacing the entire NSSEC, for example, we would be investigating that as primarily to assess if it is a transfer of control of the domain which would go through another process, a re-delegation process that involves a large amount of assessment in terms of the bona fides of the proposed new operator. So we are primarily looking for that, but a

fundamental change of all the NS records would trigger us to look closer at what is happening there.

I think it's also worth noting for this question is that we are currently developing a new authorization model for the root zone. In a nutshell, what we're trying to do there is essentially introduce a much more flexible system. I mentioned today that the admin and tech contacts that are in the root zone database are those that we cross-verify changes with.

Amongst the features that we expect the new model to have is the ability to have private parties that are just known to a TLD manager to be approvers. This will allow, for example, the CEO of a registry who doesn't want their name in the WHOIS as a contact point, nonetheless being required to approve change requests for that TLD.

We would also allow more flexibility in the sense that rather than having exactly two authorizers, you could have any number of them and that could be configured at the TLD manager's direction. So they might choose to have five people that need to authorize a change request, for example.

So essentially, what we're trying to do is build a more flexible system that's responsive to a lot of the areas of practical requests and areas where we see difficulties from time to time, address a lot of those pain points that the communities had.

And so that's something that is under development. We've floated various ideas at recent ICANN meetings, gotten a lot of feedback from our users, and we're still developing exactly how it works, but our

expectation is sort of the formal write-up of how it would work for community discussion will be probably toward the end of this year.

GEOFF HUSTON: I need to [talk] faster. I think I probably know the answer because it came up in a [noarch] meeting, but I thought I'd just ask you. There has been some exploration, different crypto algorithms, and the comment that got back was if on the ICANN side, you don't recognize that algorithm, it's unknown to you and the stuff we're using, you don't accept the DS record.

So I sort of generalize this a little bit, going, well, do you need to validate that the NSes I'm asking for, that are in the child zone as well, so I've done all the right kind of first [SNIP] test, but do you inject that they're really authoritative, that they're really delegated second [inaudible], that they're current and active?

And the same with the DS records. Do you check them all to make sure that they pass your technical test of going, "Are they functional?" And if, for example, like if you have two DS records, one works and one doesn't, you put them both in or you only put the one working? What happens if that validation fails?

KIM DAVIES: Right. So, in terms of NS records, we test all of them. They all need to be authoritative. None can be lame. We check for a variety of other things. We check to make sure that the NSSEC matches the proposed set as I mentioned earlier. We check that the glue records match the

authoritative zone for those host names. We check to make sure that the name servers are not open recursors. We check to make sure the whole set of names fits in a 512-byte response without truncating.

So there's a variety of tests. I can send you the link to the exhaustive list.

Some of those are absolutes, like you can never delegate to a lame, like privably lame delegation. Some of them, there is some discretion that can be applied. Essentially, what will happen is it will show as a failure in our system and if a TLD manager wishes to appeal that, they can come to us, give a technical justification as to why it should proceed, and we'll review that on a case-by-case basis.

Essentially, what we're looking for is not clearly that the TLD operator knows what they're doing, that it's not a regression, that as a consequence of making the change, that the outcomes will be worse, that it's made things less operable than they are today. They are the kinds of things that we have in the back of our mind when we're reviewing those requests.

And there are some things that we check for that may be anachronistic, but we do still check for, that it might be very legitimate reasons to skip. One is that we check that the serial numbers match across all the name servers, which is still a reasonable check for the vast majority of cases, but in cases where there's a registry platform that has very high frequency of updates, then there are name servers that are never coherent at any one time. So we can never measure it in a way that hits them all at once. So that's an example of something that we might to

refine our algorithm to, to have less false positives in the future. So, that's NS records.

When it comes to DS records, our absolute minimum is DS to DNS key matching, but you raise an important question, which is what happens if we can't, but it's just one of many? Essentially, we don't have a formal policy that recognizes the ability to have, let's say, standby keys that are not pre-published in the child zone. But that's in the area where if a registry operator fails a test and comes to us, and explicitly says, "This DS record is an offline key that we're not using operationally today, and we do not wish to publish it as a DS key in our zone," that's something that we'll typically approve as an exception.

We do test for validity. We do do a validation test, I should say, by trying to check the RRSIG of the SOA record of the child zone as an additional test. Now, this is not something that's published. It's something that we've evolved over the last few recent years, and it's something we expect to actually become mandatory when we next revise the set of technical checks, but it's not something that's mandatory today.

So essentially, what happens is we check for this. If we see an issue, we'll informally tell the TLD operator that we've identified this issue, and usually, if we've caught something legitimate, they'll immediately say, "Good catch. Let me fix that," or revise the change request. But it's not mandatory, so if they decide to proceed regardless, that's something that we would do.

In terms of accepting new algorithm types, for this reason, we need to be able to implement algorithms in our software. So the algorithms that

we support today are essentially crystallized from when the root zone was first signed in 2010. We have limited scope to revise that list while we're under the NTIA contract, so it was in a form of stasis for the longest time until the end of last year. Now that that's no longer a consideration, we are looking to expand a set of algorithms that we support.

True to the nature of the way root zone change is propagated, it needs to be supported in both our systems as well as Verisign, so we've been in dialogue with app developing support for additional algorithms, and that work is active right now. Our development teams have been developing support for new algorithms in the last few months. We expect to have announcements about that later this year.

So that is active work. I will say that we're not planning to support 100% of algorithm types or digest types that are standardized today. Essentially, what we're looking at is maturity of software implementations. We need mature software implementations that we can rely upon in our various tools and systems, and that's the primary determinant of what we're comfortable putting into our root zone management systems.

I think we take a fairly conservative approach. It is a root zone. It is a critical infrastructure. So we look at it through that lens. So that's kind of the history of algorithm types.

So I would say that despite some concerns in the public, we are active on that request. We are planning to introduce new algorithm types but that work is ongoing.

GEOFF HUSTON: Thanks, Kim. That's really helpful. And, of course, the next question is a logical follow-on from that. It's sort of, okay, so I want to change. You put in a change. It's in the root. My systems [inaudible] bit rough. Bad things happen. Is there any mechanism where you find a problem, what would you do? Like all the NSes go lame, the child zone changes keys and you're still running the DS from the old key, etc., etc. What is your process for monitoring, and I suppose, resolving some of these issues around the match between the zone pairing and these children zones?

KIM DAVIES: So in short, we don't do formal regular monitoring of that kind of thing outside the scope of the change request. When we receive a change request, there's a proposal to change it to X and we review the nature of the change, and we assess all the things I just talked about at that time. But we do not, on a recurring basis, independently just check these things just for data quality/health purposes.

And the reason for that is relatively simple. It's that when we last contemplated this, and this was probably ten years ago now, and there was a strong sensitivity that the IANA function should be purely reactive and not proactive.

Now the history there, and I think ICANN has evolved significantly since that time, but regardless, there was this notion that IANA should respond to change requests that are directed to it by authorized parties, but should not actively go out and start proposing changes. And we kind of, upon consideration, thought that actively sort of looking for these

kinds of problems and triggering a process, whatever that process might be, would probably fall on the other side of the line there.

Now this activity might no longer exist and, in fact, it could well be broad support for that becoming something that we do as part of our operational responsibilities. But that's not something we've really delved into recently.

I will say that also, there was also a perception that if we were to do some kind of reporting that was just available to everyone, that we detected these [slight lane] delegations, for example, that we'd be naming and shaming TLD operators that would also be not received well. So that's another reason why we've not done that kind of thing. Obviously, there's third parties that do that kind of monitoring informally, so we know that that capability is out there. It's just not something that we're doing directly ourselves.

We have considered this, but I just thought that we –

GEOFF HUSTON: You mentioned, excuse me, you mentioned that with the case of, say, the protocol crypto algorithms, you had plans, there were discussions, things were underway. Is this something that was decided ten years ago and you've just run with it, or is there any active area or consideration going on within ICANN, or IANA, or PTI on this matter, or are you relying on, say, the policy communities to give you guidance?

KIM DAVIES: On this matter, there's nothing active. I think in terms of planning, one thing that is still sort of at a concept stage is the idea of doing a new public consultation on the technical checks more broadly. So the technical checks that we have today for the root zone that I summarized earlier, that was a result of a public consultation we did around 2007 if I'm recalling correctly, so it has been ten years.

And given the relative constraints that we had under the previous IANA contract are no longer applicable, and given it's been ten years, we, as staff, thought now would be a good time to just re-evaluate all of them. So it's really pending availability of staff resources, and more directly actually, myself, we would like to recap what we're doing today publicly and ask for community input on how we might refine our technical checks.

And we were looking at it more through the perspective of tangibly, exactly what checks that we do to evaluate change requests for eligibility, but I think it could be certainly within scope to tackle this particular question as to whether it would be appropriate for IANA to more broadly monitor this on an ongoing basis and then assuming that we identify issues, what we should do with them. I think there's any number of things that could happen there and warrant a discussion, and also say that this new technology is like CDS and CDNS key records that are not quite the same thing, but would require active monitoring and we might want to implement techniques such as those as well by a similar mechanism.

So in sum, I think we have ambitions to consult with the community on technical checks in general. It's not calendared at the moment, just due

to it's one of many things we'd like to do and absent specific direction that this is a priority or what have you, we just have it on the list and we're hoping to get to it at some point in the future.

GEOFF HUSTON: Thanks, Kim. And my final question sort of stems from some sideways work that I was doing where I was analyzing the behavior of the various root server letters and those are some very subtle variations in behavior. And I sort of wondered to myself, "Does anyone check [all of] these – what is it? – I suppose thousands or possibly tens of thousands of instances of all these anycast constellations? Do they all actually publish exactly the same root zone content?

Now, do you think this is a staff matter, a community matter, an RSAC matter? You know, is anyone looking at the results of all your careful root server work?

KIM DAVIES: Not to my knowledge. It's a very good question. I mean, if I had the capability to do that, I feel as the publisher – I mean, technically, I'm not the publisher but for the sake or argument, let's assume that we want to make sure that what we've asked to be published is what's published. I would love to have a way to confidently reconcile that and know that that is happening. I think the current relationship is where we're validating that what Verisign is pushing to the root servers is correct by virtue of them publishing us the actual zone file itself that they're then putting into their distribution master and then shipping off.

But I kind of feel, from a roles perspective, whatever processes are required for root servers to validate amongst themselves is probably on the RSSAC side of the fence. But I do know today, to my knowledge, even if we wanted to and we agreed it was in scope that we have no way of even knowing all those no's are and have no way of measuring that even if we wanted to. So we're limited in our capabilities there and it could well be that the community feels it's totally outside of our scope as well.

But I think operationally, I mean, we want to be as operationally excellent as possible, and if we go a step beyond simply confirming that Verisign has pushed the right thing to going to being confident that it is widely disseminated amongst the root servers as well, if not 100%, then I think that's an extra level of confidence we have in conducting our roll.

GEOFF HUSTON: Time is moving right on. This has been really useful. Kim, thank you.

The last three questions are about the root zone KSK roll. I am not sure to what extent you and your team are involved here, but I'm asking around not the process. I think we're all familiar with that and that's easy in terms of it's not easy in practice, but understood.

But I was actually touching upon this issue of what if you can't get to both keys. In other words, both locations collapse and eat the bricks at the same time. What if, for some reason, you believe it's being compromised? What if there is a need, however it came through, to say, "We need to roll this key and we need to do it now"? Do you have any

procedures in place around use of this key outside schedule, Windows publication of the new key, outside scheduled events?

KIM DAVIES: Right. So we have discussed these issues. I mean, I guess to your first question, are there procedures in place to manage an emergency KSK key roll? PTI and Verisign have been collaborating on this topic. We've had a series of meetings in the past year in terms of disaster of planning, both for an emergency KSK roll, as well as an emergency ZSK roll.

It's worth noting that if Verisign's facilities had a problem, they needed to roll the ZSK. That would also require an emergency key ceremony as well, even if the KSK itself was perfectly fine.

So we've produced some internal documentation in terms of organizations in terms of staff response, what the game plan is, in terms of how quickly we're expected to respond, fire what mechanisms, and so forth.

The KSK management staff reported to me until about three months ago. They've transitioned to a new Director of Security that we recently hired. So for that reason, I'm not exactly sure the exact state of those, but that work is still ongoing. We do have some plans, but I think that we want to continue to expand upon them.

The second question, are there procedures in place in the unlikely event that two KSK repositories are inaccessible? So, it really depends on what inaccessibility means there. We only during regular affairs, need to

access the KSKs every three months. So there was some intermittent inaccessibility that could be recovered from. It's entirely reasonable that we might need to delay regular operations, but it wouldn't fundamentally has to do anything emergency-related once we've established that the facility remained secure throughout the whole incident, whatever it might be.

But, you know, worst case scenario, let's assume that both KSK facilities are obliterated by some event. In that case, it's relatively straightforward that we need to regenerate the KSK from scratch, and I think in that case, we would need to reconstruct key management facilities from scratch as well.

I don't think we have an exact procedure for how that would happen. I think, firstly, it's quite far-fetched, but also that it is, it will be dictated by the circumstances of the event while in large part as to how that would happen.

So it's probably something we need to do more about, but we don't have an exact game plan as to how we would reconstruct key management facilities from scratch should both KSK facilities be destroyed.

And then other –

GEOFF HUSTON: Is this a topic that we should directly take up with your new Director of Security?

KIM DAVIES: I mean, it would be an area for them to consider. I'm not sure that they've actively thought about this topic more than I have, so at this stage, I'm not sure that [inaudible] much. But this could be an area, certainly, in your recommendations that you might want to direct us or encourage us to consider more fully. I'm not quite sure what the right answer is there, but I mean, I think just from a logical perspective… I mean, the KSK is only held in those two facilities. We have backups, but the backups are against things like hardware failure, so the backups themselves are in those facilities as well, I think.

GEOFF HUSTON: So in the case when the [inaudible] died on one of the [XSMs], there was enough [inaudible] to get you back.

KIM DAVIES: Yeah, I mean, we have duplicates of the machines themselves, so battery failure in itself shouldn't be an issue, but let's, hypotheticals we've talked about as some kind of bad firmware issue that they all have the same firmware, so all the HSMs might act radically or erase their content in a similar fashion.

We've tried to, as much as reasonable, accommodate all the single points of failures we could conceive, and battery is one. So we made sure that all the HSMs have different batteries from different batches in them. But we had to settle on a single vendor for the HSMs themselves. We don't have an alternative vendor that is mutually compatible with these HSMs that we could store it across different models.

So there's some natural single points of failure, and sort of the backup regimen is really designed to mitigate and start even as far-fetched as those kinds of scenarios might be.

But the back-up [inaudible], we don't have a third facility, let's say, that's at the same standard as these two KMFs to store the backups in, so we don't have an alternative place to put those backups that is resistant to physical threats that might destroy both KMFs. So that is a tradeoff and that was the original design of the KMFs. There might be something worth revisiting, but that's the way it is today.

GEOFF HUSTON: Okay. We're right to the top of the hour and I'd, again, just like to thank you, Kim. This has been extraordinarily valuable to me and for the others who have been listening in. I think it's sort of been or come to be the body of a lot of what I'm going to report back on, so again, I'd just like to thank you for your detailed and informed responses. They are most helpful. Thanks, indeed.

A bit just on housekeeping, I have no plans for a meeting next week, or indeed, any further meetings until I get clearer guidance from the Chairs, Eric and Denise who came in and out, just to the help of this group and level of participation.

My own plan is to take these responses from Kim and draft up a more substantive report, including some recommendations and circulate that back in around two or three weeks so that by the time we have a face-to-face meeting at the next ICANN meeting, we will have a substantive

part of this work done from this subgroup, and hopefully then, we sort of can see the end of this particular subgroup work.

So at that stage, things are looking good, at least from my perspective, that we're kind of past the investigation stage and now into the write-up, which as I said, I'm not sure I may continue with meetings for. So I will not be requiring a meeting next week, and it would be very good if Denise would want to add some more guidance, I'll happily accept it. But that's my current plan.

Okay, well, thank you, Kim, and Eric and Don for joining us. Much appreciated. And thanks to Karen and Yvette and Bernie.

Won't see you next week, but we will see you each other at the ICANN meeting, if not before.

UNIDENTIFIED MALE: Thank you very much.

UNIDENTIFIED MALE: Thanks a lot, Geoff.

UNIDENTIFIED MALE: Thanks, Geoff.

UNIDENTIFIED FEMALE: Thank you.

GEOFF HUSTON: Thanks. Bye.

**[END OF TRANSCRIPTION]**