
MONTREAL – RSSAC Work Session 2
Sunday, November 3, 2019 – 10:30 to 12:00 EDT
ICANN66 | Montréal, Canada

ANDREW MCCONACHIE: Hope everyone had a nice break. This session, we have an hour to talk about RSSAC002v4 and then we have half an hour where we're going to be getting a presentation from some academics who did a legitimacy study of ICANN.

Let's start on 002V4. Ozan, can you advance the slides? Yeah, perfect. The current version is V3, this is Advisory and Measurement to the Root Server System; thank you, Ken Renard, for stepping up, for helping to guide this effort, thanks a bunch. The process is going to be pretty much the same as it was, you know, what we just did for 023 and 026.

Staff, me and Danielle, have made changes to 002V3 and to a draft 002V4, and first I'd like to go through the various topics that were brought up on the list and see if people have other topics that they want to add to that list? The original list of topics -- or could you advance one slide, Ozan? So this is the original list of topics -- or I shouldn't say original -- this is a list of topics that came out of discussions on the caucus mailing list a few weeks ago; I think like six weeks ago is when the discussions started.

If you've been paying attention to the list, you've probably seen this list of topics before, but I will go through them again. The first one is

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

something else that came out of the Metrics Work Party with regards to latency and staleness, trying to differentiate between latency of a query and its response and latency of publishing a new zone. The edits I've made to 002V4, hopefully clarify that. Actually, when I made those edits, there were fewer things I had to change than I originally thought I was going to have to change. The 002V3, I already did a pretty good job of that and I think I only had to make a couple changes.

Number two is measuring zone size, where that should be done. The size of the zone measurements has been moved into a new section I believe. Folks can interrupt me at any time if they have comments on these.

Item number three, there was a suggestion on the list to provide some kind of QNAME minimization measurement. Based on discussion on the list -- and this is a topic for right now -- but based on discussion on the list I did not include draft text on that because I didn't think that -- what I read from the list was that more people against that measure than actually for it. We can go into that. Does anyone want to say anything about that? Does anybody strongly believe that RSSAC002v4 needs to do something for QNAME minimization measuring? Duane.

DUANE WESSELS:

My interpretation of the current introduction is that QNAME would be out of scope for this document.

ANDREW MCCONACHIE: I'm hearing that we just -- that won't be followed up on and we won't be doing anything for measuring QNAME minimization id 002V4. Moving on to 4. I guess there was a change between V2 and V3 regarding how to count broken DNS queries, and there was some confusing and, Duane, you sent me a Github link which had a whole bunch of discussion on some confusion around how to implement counting different types of broken DNS packets, so I added some text that hopefully clarifies how to count these. We can have a greater discussion on that, but did you want to say anything about that?

DUANE WESSELS: No, I think you captured it. That was from Jerry at DNS OARC who had implemented against this document and had some questions.

ANDREW MCCONACHIE: Okay. Number 5 is removing quotes around start period, just the term. I don't know why start period got single quotes but it had single quotes and there was an ask to remove them, so I removed them.

Number 6, also in that Github link that Duane sent me, there was very little discussion in 002V3 about how to represent null and zero values. I went and looked at the YAML specification to try to get a better idea. YAML actually has multiple ways to represent null, so I think in the document I just said, "Do what YAML does." Then there are some explicit changes in the document stating which measurements need to record zero and which measurements do not need to be included and things like that. So, we can go through those changes. There was

another suggestion about including concatenating multiple YAML documents into single files, that didn't seem very controversial on the list, so I made -- Paul, go ahead.

PAUL HOFFMAN: You can't do that in YAML. You can't just concatenate files together, you'd have to make a list. You'd have to make a super list and then concatenate them together. If you just stick two things together it's invalid in YAML.

ANDREW MCCONACHIE: Okay, that's good to know.

PAUL HOFFMAN: Sorry, I have done that painfully.

ANDREW MCCONACHIE: That's good feedback. There was a suggestion to expand the URL standard specified in Section 5.7, expand it all the way to the left so the only thing that would be unique would be the domain name. We could go over that. I get some comments on this one. I saw Michael first and then Brad.

MICHAEL CASADEVALL: I actually had a comment on the null and zero values, so you can come back to me.

ANDREW MCCONACHIE: Brad, did you have a comment on this URL?

BRAD VERD: No, my comment on was on number seven.

ANDREW MCCONACHIE: Michael, go ahead.

MICHAEL CASADEVALL: As far as null and zero go, I think we need to explicitly define these terms and not look at the YAML specification because I did look at that. The YAML specification essentially says it maps to null data types in various programming languages. The problem is various programming languages have interesting definitions of what null means and all disagree with each other; anyone who's played with SQL databases know exactly where I'm coming from. I think we need explicitly define what a null is and zero is, just so there is no ambiguity in our documents.

ANDREW MCCONACHIE: The document doesn't go that far, the document just says, in order to represent a null value, use these strings. I think it's a dash is one of them, you can use lower case n-u-l-l, upper case N-U-L-L, there's no guidance in the document about how to deal with those values, how

to shove them into a database, that's kind of outside the scope of the document.

MICHAEL CASADEVALL: Null is the actual datatype in YAML though; it's not a string or other comparison. You can have a string that says null and then you have the actual null type in YAML files. This is the ambiguity I was pointing at.

ANDREW MCCONACHIE: Duane, go ahead.

DUANE WESSELS: I believe the concern -- this came from Jerry, from his implementation experience. I had the same confusion with him when I was talking about this with him, but I think his question was not so much like how do you represent the value zero or null, but he had cases where we have these lists; if the list is empty, how do you represent that? We can go back and look at the specific case if you want but just to keep in mind that I think that's where his confusion lied, not so much in do you write down n-u-l-l or zero or anything like that. There are many ways that things can be null, lists can be empty, arrays can be empty, values can be zero, maybe all these things need to be looked at.

ANDREW MCCONACHIE: Okay. Did you want to add to that, Paul?

PAUL HOFFMAN: I was going to agree on that because lists being empty is defined quite easily and null values, as Michael said, in YAML are hard and when I looked through this, it was not clear if I'm supposed to be reporting something and I don't have a value, am I supposed to put the name and put a null value or am I supposed to elide? I think that that can be done and I wouldn't be surprised if that was Jerry's problem because it has been mine, of that we just have to go through everything where there might be -- you might not have a value and say what to do if you don't have a value there.

ANDREW MCCONACHIE: Michael.

MICHAEL CASADEVALL: Generally, the way I would handle this, coming from a database perspective, is null should be used for information is unknown, if it's known that we don't have value X, we either don't include it or we have a cardinal value for it. But we need to define how this works or we are going to cause ourselves pain and misery.

ANDREW MCCONACHIE: I think we need to drill into the document on that one. But first, Brad, you had a comment on the 7.

BRAD VERD: Maybe this goes to Duane because you're kind of channeling Jerry. Because the answer to this one was, you can't just concatenate YAML files but what were they trying to do and maybe we need to address it in a different way. I'm just not sure what the person was trying to do when this issue came up. I feel bad saying, sorry you can't do that, I'd like to know what they were trying to accomplish.

DUANE WESSELS: I think Jerry's frustration was that it just ends up being a lot of files on the disc the way it's currently specified, you just output a lot of small files and he was thinking, maybe it's more efficient if his software at every interval could output just one file with all the things together. To fix it, we would have to redefine the format, which would be kind of icky, I think we just live with it as it is, even though it's a little bit unfortunate.

ANDREW MCCONACHIE: Okay. The final item on this list is looking into recommendation three on RSSAC002v3. Recommendation three is the recommendation that says, the RSSAC will look at this again in two years and determine if it needs to be updated. Just like with I think RSSAC023, we should figure out what the guidance is in the document for how often this document get's updated or is there a need for guidance in RSSAC002 about how often RSSAC002 should be updated or is it just something that the RSSAC doesn't want to specify? Any comments on that, Liman?

LARS-JOHAN LIMAN: I suggest that we don't specify a specific timeframe, we suggest that it should be reviewed from time to time. Tying a whip to our own back for doing -- if we keep doing this in many documents, we will do nothing else but review our old documents in five years down the line. I suggest we don't put a specific timeframe in there.

ANDREW MCCONACHIE: I think the strongman text I wrote says, as necessary, that's to be interpreted. Any other comments on that?

DUANE WESSELS: Not directly to that but related and I see you've highlighted this in the document, there's a metric for number of sources, unique sources and it still says that this is an optional thing to collect and that it should be reevaluated after three years, I guess we need to think about if we want to have it as optional or if we've all decided that by now it should be there and change that text or not?

ANDREW MCCONACHIE: That's the 10th item on that list then, is what you're saying? Evaluate what a unique source is, needs to be in the document at all?

DUANE WESSELS: Yeah, I guess so. You've highlighted it already in the draft but I guess it does need to be number 10 on the list.

ANDREW MCCONACHIE: Are there any other additions to this list that need to be made? Hearing none, let's go look at the document.

Ozan, could you scroll down to Section 3. I'm sorry, not the statement of work the RSSAC002v4, Google Doc. Do you need the link? I just sent it to you. If you could post that link in chat as well so folks can follow along at home. The people who do have the document loaded, I want to start in Section 3, RSO Measurement Parameters. What I'd like to highlight here is the size of the overall root zone has been moved to a new section on Capturing the RZM Measurements.

I think this has always been, correct me if I'm wrong but I think this has always been something that the RZM does measure but we never had a separate section in the document for RZM Measurement versus RSO Measurement. I think the suggestion was to breakout this measurement to a new section, which is now Section 4. Does anyone have any comment on that? It doesn't change any of the text, it just adds a new section. I will take that as noncontroversial.

In Section 3.1, in the first paragraph there, you'll see one of the changes I made which relates back to Item 1 on the slide, which is clarifying the difference between publication latency and query latency. In a few places in the document I added the publication and there were different words used, sometimes there wouldn't be a word like distribution news, sometimes it would just be talking about latency without specifying what kind of latency. I normalized it everywhere to be publication latency because I don't think this document talks anywhere about query latency. I'll will move on.

I think the next thing to talk about would be in Section 3.5 on Numb Sources and, Duane, this gets back to your addition to the list. At the bottom of that page, the first page with Section 3.5 Numb Sources IPv6 Aggregate, there's a paragraph there that I've highlighted, which say, this set of metrics is marked as optional for a three-year period following the acceptance and publication of Version 1 of this document by RSSAC. The question to the group is, do we want to keep this metric or do we want to remove it?

DUANE WESSELS: My opinion would be, keep it, but either keep it as mandatory or keep it as optional.

ANDREW MCCONACHIE: Liman?

LARS-JOHAN LIMAN: I'm not totally averse to keeping it, but has it ever been used or are we just collecting statistics that we think might be used in some context that we don't know in the future?

WES HARDAKER: I don't know positively but I'm pretty sure I can speak that I'm colleague has used it and the reality is, we have no idea, there's no way to measure whether it's been used. That being said, I'm actually fairly positive that my colleague has probably used that data.

ANDREW MCCONACHIE: Michael?

MICHAEL CASADEVALL: I would seriously consider removing numb sources IPv6; the reason for this is due to IPv6 privacy extensions causes clients to renegotiate the bottom 64 bits of their address at regular intervals. That will cause false positives on a regular basis because a device, I monitor this, I've seen devices change IP's 60 times within an hour. That metric in and of itself may not be useful and it also bloats up the statistical information considerably, keeping the aggregate below it, is much more useful because slash 64 is the smallest possible routable address and it voids all the problems of looking at individual IPv6 addresses.

ANDREW MCCONACHIE: Any comments? The suggestion is to, as I understand it, first off remove numb sources IPV and then keep the other numb sources as optional. Any comments on that? Paul?

PAUL HOFFMAN: This really goes back to the purpose of 002. If the purpose is to look at the scaling of the root zone, I'm not convinced at all that these numbers are valuable for that. I'm not questioning that someone can't use these numbers because standardly in diddle data, whenever you see an OARC presentation that uses diddle data, they use this, they don't go to the 002.

If we are meant to be collecting things that are useful to researchers than this can't be optional because it may be that a root server operator stops doing it in the middle of a week or whatever. It will not be useful if it's optional. I'm not convinced it's useful, so I'm saying we need to keep it but I think having optional ones makes the entire thing unuseful to researchers.

DUANE WESSELS:

Regarding Michael's suggestion about removing the V6 addresses, to me the only reason to remove it would be if it became burdensome to compute and I don't think that's the case right now, mainly because V6 usage's is quite a bit lower the V4 usage. As far as I know, 12 of the 13 root server identities are publishing this data already. I think it's doable. I guess I don't see the reason to remove it, to remove that one sub metric.

ANDREW MCCONACHIE:

I saw Michael and then, Paul, did you have your hand raised? Okay, just Michael.

MICHAEL CASADEVALL:

The reason for removing it is it's not a useful metric and it is incredibly misleading because of privacy extensions. If I really wanted to, I could hit a root server with 64 -- two the power of 64 IP addresses and completely skew the metrics entirely because you cannot tell how many individual devices are coming with a given IPv6 address, if there

is one, if there is multiple, so forth and so on. Privacy extensions are enabled by default in Windows, Android, IOS and Mac OS10 and some Linux distributions. It's at best misleading and I can't even think of a use for that information if you understand how IPv6 actually works in practice.

ANDREW MCCONACHIE: Wes?

WES HARDAKER: I actually like Michael's suggestion; when he first made it, it made a lot of sense. Whether or not it's statistically interesting, I certainly understand how the privacy extension works, even though there is debate about whether it actually gives you privacy, that's a whole other debate. It seems to me like putting it in is functionally and venerability to a large extent because we don't know, nobody's done exactly what Michael just described, which is send us a packet from a whole bunch of addresses and watch us try and calculate it.

I think we might be able deal with networks, although whether V6 and BCP38 is widely deployed, it could be bigger than two to the 60, there is that issue too. It seems to me like putting that in and saying that we should be doing that calculation is simply dangerous, where as if we go to the slash 64, the slash 48, that's a safer thing to do and is probably equally as helpful.

ANDREW MCCONACHIE: Am I reading the room, is to I guess remove Numb Sources IPv6? There is another question about whether the remaining metrics should be optional or required? To Paul's point, are these metrics useful if they're only optional? Is everyone okay with these metrics just being optional? Go ahead.

DUANE WESSELS: I think they should be made nonoptional.

ANDREW MCCONACHIE: You mean required?

DUANE WESSELS: Yeah, just like all the others in the document. This is the only one that's marked as optional and I think we should take that out.

ANDREW MCCONACHIE: Okay.

WES HARDAKER: Operators that consider it a burden should speak now or forever hold their peace.

ANDREW MCCONACHIE: That's kind of what I'm waiting for and no one seems to be speaking. Then my reading of the room is that we remove Numb Sources IPv6?

In fact, I'll just make the edit right now. Also, remove this paragraph here, seeing as we've discussed this, and where do we discuss them being optional? I guess it was in the paragraph I just deleted.

DANIEL MIGAULT: Just a comment. I thought every three years we'll revise the document?

ANDREW MCCONACHIE: That's what V3 says right now, it says every two years the RSSAC should look at this.

DANIEL MIGAULT: Okay, but you've removed the --

ANDREW MCCONACHIE: No, that's just for this specific metric. These three metrics had a specific paragraph saying, these are optional from V1 and please look at this later.

Moving down, I mentioned this earlier, there's a new Section 4 on RZM Measurement Parameters, previously this was just in the RSO Measurement Parameters Section, which was kind of confusing because it's not actually something that RSO's do. Any comments on that? Having seeing none before and seeing none now.

Let's skip down to Section 6. The first sentence in Section 6 says, metrics should be stored in per day, per metric YAML formatted files. I note here in this comment that YAML has multiple versions and this document doesn't specify a version of YAML to support. Is there interest in specifying a YAML version? Michael?

MICHAEL CASADEVALL: Before I answer that, was there a specific reason why YAML was chosen over JSON? Because that's going to influence my answer.

ANDREW MCCONACHIE: I don't know the answer to that.

MICHAEL CASADEVALL: The reason I bring up that particular point is that a JSON formatted document is valid YAML and a YAML parser is very heavy and very [inaudible], and a lot of data can lag down. I've seen this in production systems where switching from YAML to JSON gives you a fairly large performance increase. We could actually say the data is in JSON and not break anything due to the superset status of YAML. I'm questioning if that's something we want to discuss?

ANDREW MCCONACHIE: Daniel?

DANIEL MIGAULT: At the time the document was edited, I raised the same question, knowing nothing about YAML nor JSON, and I've been told that from what I remember, that YAML was more efficient.

PAUL HOFFMAN: It's not. At the time or I don't know at the time of this but it was considered for a while to be more efficient, it is more efficient for writing and it is vastly less efficient for parson, they have been even competitions to make more efficient YAML parsers and they still aren't as good as JSON.

Having said that, there is no reason to change it now. I am not convinced that anyone who is reading this data and parsing it, is doing so very often or is under any time constraints and therefore the cost of making a massive change to be more efficient would I think totally be lost and once we start doing that, then well bike shed should be JSON or in fact should it be CBOR of whatever. I think we can just drop that one.

As for you question though on a specific version, yes you should name a version because the folks who are doing YAML are having the same problems that the folks who are doing JavaScript have, which is they see problems that they created a few years and they desperately want to fix them but they don't want backwards compatibility but the people who don't want to break backwards compatibility are older and therefore age out and there is a reasonable chance that YAML could have a breaking change in the future. They would seem really

minor to them but would then screw us. I would say, pick a version that we know we are compatible with and just say that's what it is and even 30 years from now, that'll be just fine.

ANDREW MCCONACHIE: Michael?

MICHAEL CASADEVALL: Going back to the original question, I would then recommend we standardize in the YAML 1.2 specification, that's a strict superset of JSON and is what I believe we are currently using de facto.

ANDREW MCCONACHIE: I really don't know enough to say whether or not every RSO currently produces valid YAML 1.2 files. I guess I'm just a little bit concerned that we say 1.2 in this document and there is a whole lot of work for something we didn't think of that causes some major compatibility issue. Paul, go ahead.

PAUL HOFFMAN: I will volunteer, not immediately but over the next few months, to actually analyze the files with a strict 1.2 parser. I would be surprised if anyone is not omitting 1.2 because the differences between 1.1 and 1.2 were small. My concern is differences for future versions of YAML. If that's a concern of yours, I will find a strict error reporting parser and try to run it through the corpus that we have or at least recent ones

and just do that and if that's fine, then I agree, we can just say 1.2 and be done with it.

ANDREW MCCONACHIE: But even if it fails sometimes but the changes are minor, that would be really good to know. Duane, go ahead.

DUANE WESSELS: I have some experience in analyzing this data. In my experience we have more cases where it's not a valid YAML file in any sense rather than YAML 1.2 or 1.whatever. We already have to deal with this. Now, of course going forward we should be as good as possible, we should make it as clean as possible but we already have to deal with this.

ANDREW MCCONACHIE: Thanks Paul, for doing that work. That's an action to Paul for -- I'm reading the room that we tentatively specify YAML 1.2 here and that Paul's going to do some research to make sure that actually works on recent current files. Go ahead, Paul.

PAUL HOFFMAN: I wasn't going to bring up with Duane just brought up but, he's right, that going through the documents you can easily find broken YAML and you can see what they meant because YAML is text based. How does this group want to deal with that in this version, if at all? That's is, if someone emits badly formatted data, do we want to mention that

in this document? Do we want to make requirements? We can't make requirements because this is somewhat voluntary but is that something we want to deal with?

Again, I wasn't going to bring it up but since Duane did, I'll follow along. It happens a reasonable amount and YAML is hard to do quite frankly. For those of you who aren't aware, ISC is adding YAML to Dig, the next release version of Dig you'll be able to say, Dig plus YAML and get YAML output and they've had hard time getting correct YAML every time as well. I've done enough bug reports, is that something we want to deal with here or not? The fact that some files are not meeting this thing that says, should be.

ANDREW MCCONACHIE: Go ahead, Michael.

MICHAEL CASADEVALL: The data is pretty useless if we're spitting out malformed data that requires someone to go through and hand edit and due to the way that YAML has context sensitive whitespace, that is not always the easiest thing to do, which goes back to the original discussion, do we actually want to change out the format? I can already see the heads shaking. I just feel like, we need to standardize and if we're having trouble getting these files in the right format, then a standardized library or something that helps, is going to be essential but it has to be one format that can be machine parsed or we're basically making a lot

of information that no one can use without beating their heads against a rock.

ANDREW MCCONACHIE: Wes?

WES HARDAKER: I suspect that the only thing that we can do with respect to compliance is putting in a note saying, please contact the operator in question if you discover that some files are not parsed and hopefully, they can fix them. I mean, certainly using a library to produce the YAML is a safer way to do it. Things do go wrong, YAML's not as easy as it looks, that's 100% true. It's easier to read and researchers do actually just look at raw data than JSON for example but I'd just put a note. There is no way -- we can't put in punishments or anything else, what else are you going to do other than, please contact the operator if you notice broken YAML files and have them fix it.

ANDREW MCCONACHIE: Duane?

DUANE WESSELS: One thing we could do is, this document could say that if invalid YAML is discovered, then the operator or the publisher should or is allowed to go back and edit those. I think for the most part we tend to be very hands off about revising older files, we don't want to do that. It looks

like revising history but we could say that that is an okay thing to do, to fix formatting problems.

WES HARDAKER: I think it actually has been done to be honest. But no, they shouldn't change the data, you're right.

ANDREW MCCONACHIE: Daniel?

DANIEL MIGAULT: I have two questions. The one is, the things that we observe and require some manual checks, would JSON for example solve that problem or no? No. Okay.

ANDREW MCCONACHIE: Paul. Go ahead.

PAUL HOFFMAN: It's easier to create broken YAML than it is to create broken JSON. It is perfectly possible to create broken JSON as well. The only advantage here would be that there is exactly one standard for JSON and that as it gets changed by the ECMA, we can still refer back to the RFC that specifies it. I don't think it's a big win here, quite frankly if people are willing to back and fix their YAML, which they would because quite frankly almost all of the breakage that I've seen has been

mechanically reproducible as in, oh they added a space here. I think that that get us the benefit that we want.

ANDREW MCCONACHIE: What I'm hearing is that we may want to add some language to this document, helping consumers of this data, telling consumers of this data that if they encounter broken YAML, they should contact the RSO. That's an action item to Staff to draft that language. Brad, did you have something on that?

BRAD VERD: I was going to say we should compel the RSO's to produce proper files and fix ones that are broken.

ANDREW MCCONACHIE: That's not something for this document but yes. Moving on. Down at the bottom of the page, on one of the bullets. This text here, some keys must -- this is new text that I added, some keys must be included, even if their value is zero, while others must not be included if their value is zero. Values of zero must be represented by the number zero, Unicode compliant 0030.

Then there's an additional bullet, no value representation must use the YAML 1.2 specification. Look, I actually say 1.2 there, sorry about that. Kind of jumped the gun there. Must follow the YAML 1.2 specification of representing no values. Comments on these two bullets? This is just about representation. Paul?

PAUL HOFFMAN: This isn't a comment just on these two bullets but it's a comment on the use of the word null, which appears in Section 5.1, which is right below this. I don't know if you want to wait for a moment? I figured you were going to skip over 5.1 but it relates to the null part of this.

ANDREW MCCONACHIE: Okay. We can definitely talk about that. I guess I'm interested in -- does anyone have a problem with this language specifically about how to represent these values? Michael?

MICHAEL CASADEVALL: I do because the problem with the language as is, is that there is no way to define that information was unknown for a given time period. Statistic system goofs and a bunch of information is lost, so now we don't have a way of saying, we already deleted it from there, we don't have a value that indicates this is unknown, which is what the typical definition of null is verse information is present versus not present. It's trinary logic, which can be a little hard to get your head around.

ANDREW MCCONACHIE: The difference between unknown and not captured? We're collapsing both of those values into null and that's your issue?

MICHAEL CASADEVALL: More specifically so, you can have -- during the reporting period, load time information was lost due to system glitch; that should be null because the information is unknown, it's not that it's not present, it's simply not available or known, that would be the definition of null. If the information was not reported, then it should be a blank value or zero or a cardinal value as defined or the information will be present. Hoping I'm making this point clearly.

ANDREW MCCONACHIE: So, how would you -- Paul, go ahead.

PAUL HOFFMAN: Except we only have one place where null values are even mentioned and that doesn't count there but it is exactly the question I have in Section 5.1. The word null doesn't appear currently anywhere else in the document.

ANDREW MCCONACHIE: So, let's skip down to 5.1 and I believe you're talking about that last sentence there Paul, of Section 5.1, if no load time metric is available, it should be represented as null. Previously that had been written has, it should be marked with dash. Go ahead, Paul.

PAUL HOFFMAN: Two things that are not clear to me in this, if the load time metric is not available or no load time metric is available, is this even produced

at all? Is what we mean is, that this metric is not produced at all or that the thing that appears after the time and the date stamp is what you're calling null, which we could call zero? And I don't know what the intention here, does that make sense?

I don't know if anyone ever has this situation where they have no load time metric but to me, the original was saying, produce this but produce it indicating that we don't have a load time metric value but the way this is saying, it should be produced as null, we could mean, you don't have this section at all. I would defer to the people who are creating data here of, have you always been producing it in the past a dash and if so, then we could easily just say, zero or minus one. As Michael had said earlier, we could just have a sigil that makes no sense as a real value and put it there.

MICHAEL CASADEVALL: Off the top of my head I don't know, I'd have to go look at the data.

WES HARDAKER: I think I can speak for us, that there are times where -- I think there are a couple instances where we lost RSSAC data due system failures that weren't necessarily visible to the rest of the internet but it was the internal processing data or something that caused us to, we should have put in nulls, I think we did in anything but I'd have to go back and look.

PAUL HOFFMAN: If that's the case, I think instead of saying, should be represented as a null value, we could just say, represented as minus one, which is an impossible time, very clear to find, things like that.

MICHAEL CASADEVALL: I will agree with Paul,;the only time we should ever see null in the metrics is if something went horribly wrong and data was lost, under no situation should be using null as a place to represent, we are not collection this data, we should have a cardinal value of like negative one.

ANDREW MCCONACHIE: In that case, we could just keep the text as it was in V3, marked with a dash or would you like to use a different sigil than just dash? Paul, go ahead.

PAUL HOFFMAN: Absolutely, I would like to have a numeric sigil so that the parser -- normally, YAML parsers and JSON parsers want to know the type that will appear as a value and a dash would be a text type as compared to a numeric type, if we use minus one here, then it's numeric types.

ANDREW MCCONACHIE: An annoying question for someone who really hasn't written too many parsers, if you're trying to store a negative one into an unsigned integer...

PAUL HOFFMAN: Don't do that.

ANDREW MCCONACHIE: Okay. I mean there is no other sign for all these other values, I worry about someone just using unsigned integers everywhere and then they're like, negative one, it's this giant value.

PAUL HOFFMAN: As far as I know, there are no YAML parsers that do that because YAML doesn't have a concept of unsigned for resulting values. This was a problem we hit when we developed CBOR.

ANDREW MCCONACHIE: Okay. Let me make some changes to this text. I'm going to back out the changes I made and simply make that change. Michael, go ahead.

MICHAEL CASADEVALL: As [inaudible] you're saying you put a string saying negative one, it should be represented as the integer value negative one to completely ambiguous.

ANDREW MCCONACHIE: It should be represented as the integer value of negative one.

MICHAEL CASADEVALL: Perfect.

PAUL HOFFMAN: I'm sorry, I'm going to be picky here, don't put the period at the end of that sentence please.

ANDREW MCCONACHIE: Fair enough. Given that this is the only place in the document where the word null appears, other than the directions in Section 5, this sentence and the immediately preceding section, where it explains how to represent null value, I guess we no longer need that sentence. Paul, go ahead.

PAUL HOFFMAN: I'm not even sure why you need it here because if you need it here, if the load time metric is not available, aren't any of the metrics possibly not available? Why is this metric called out as one that's possibly not available and all of the other metrics are not called out as possible? I would imagine -- Wes brought up the case of, there was some internal fault that we had, I would imagine that could be true for any of the metrics.

ANDREW MCCONACHIE: Duane and then Wes and then Michael.

DUANE WESSELS: I think it's less about whether -- it's not about whether it's available or not but it's how to represent it. Not all the metrics have the same representation, some are lists and things like that. How do you represent lists when the data is unavailable?

ANDREW MCCONACHIE: Michael?

MICHAEL CASADEVALL: The other question and part of the point I want to bring is, do we want representation metrics that data is lost due to technical reasons, in which case is a valid place to use null. I'd be completely happy removing null from this document entirely.

ANDREW MCCONACHIE: I think the word only appears in one sentence right now, so we could just strike that sentence. I'll null it out, nullify it. Hearing no objections, I will delete that sentence. Actually Ozan, can you scroll up a little bit just so that everyone sees. I'm talking about the second to last bullet, hearing no objections, that sentence has been deleted.

Ozan, can you scroll down to Section 5.2, The Zone Size Metric. I've added a sentence here similar to the discussion we've having about what happens when operation is not normal and values aren't recorded properly. I added this sentence here, although it should never happen in normal operation, zone sizes with a zero size must be included. There is a lot of other sentences like this that I added for

each individual metrics which basically say either, the metric must be included as zero if it's not recorded or it should not be. Any opinions on this sentence? Michael.

MICHAEL CASADEVALL: I have a bigger question on this metric at all is, why is this useful on a pro root server basis, they should all be -- unless something has gone horribly wrong, it should all be running the same root zone file, which we could collect somewhere else but I don't actually think there is a point of collecting this on individual root.

ANDREW MCCONACHIE: It's not, it's for the RZM. That new section that was added where there is one, this is the single RZM metric and all the other ones are RSO metrics. Hearing no objection, that sentence will stay.

Can you go down to 5.3, Ozan? Like the last sentence we talked about, I added a sentence here, only traffic volumes with non-zero counts must be listed. Karl, go ahead.

KARL REUSS: What's the reasoning for that? It seems like it would be useful to know that nothing happened during this period on a certain transport or address.

ANDREW MCCONACHIE: I'm basing this on -- I think this was the discussion that happened on Github from Jerry, if I remember correctly why I made this change. You're saying that for all the different traffic volumes, even if they count is zero, they should be listed?

KARL REUSS: I would think they should be listed, the fact it didn't get anything but it was there and capable of receiving it seems interesting.

ANDREW MCCONACHIE: Paul.

PAUL HOFFMAN: Technically it doesn't have to because these are dictionaries, so you could have something without value. I think this brings up a constant problem we have in this document throughout which is, that we mix up the people who are creating this data which is y'all, and the people who are analyzing it, which is some of us all. But someone analyzing this, who didn't see that dictionary value would know that it had to be zero because it was a missing dictionary value, that just comes with any parser. I don't see an advantage or a disadvantage of adding this.

KARL REUSS: Gotcha, thanks.

ANDREW MCCONACHIE: What you're saying, Paul, you don't mind if the sentence were deleted or kept? Karl, do you agree with that? Okay. I guess there is -- I'm of the opinion then, if no one feels strongly that this sentence needs to exist, I just don't add it. It doesn't sound like there is much support for this sentence. I'm going to suggest deleting this sentence.

Let's stop here because we're out of time. We'll pick up again in Section 5.4, we can just note that that's where we stopped. I'll hand it back over to you, Fred.

FRED BAKER: Any place that's open. There's an open spot over there, maybe you go over there.

JAN AART SCHOLTE: Does anyone see a red dot? I don't. No, okay, that's alright.

FRED BAKER: Let me introduce our guests; we have [inaudible] and I'm probably killing your name, two professors from the University of Gothenburg, who have been doing a research study on the legitimacy of ICANN. They asked if they could share their results with us. Over to you.

JAN AART SCHOLTE: Thank you very much, Fred. Thanks everyone for coming along and listening to us this morning. The biggest thanks go to many of you who actually contributed to our work. The interviews were done

under condition of confidentiality, so I can't say thank you in name in this open space but I thank very much and I see a lot of faces that I recognize, so thanks so much, we couldn't have done any of this without your participation.

What we're doing here is giving you some initial findings about levels and patterns of legitimacy at ICANN. We're going to say a few things about the study, just underlying again it's an academic and independent study, it wasn't commissioned by ICANN or anyone else. We're going to tell you about ICANN's legitimacy as seen from outside ICANN, the wider world, the wider world elites, wider world public.

We're going to tell you about ICANN's legitimacy as seen from the inside by the Board, by the Community, by the Staff. We're going to say a few things about what your constituency, meaning the technical community, what you think are the priorities and what you think ICANN does well and not as well and compare you to the other stakeholder groups.

Just to give you a couple of headline messages before we start looking at graphs and things. A few headline messages for you to take home with you after [inaudible] of all of this. Average legitimacy beliefs towards ICANN are not high enough to be complacent and not low enough to be alarmed. They are kind of in the middle. The overall verdict is moderate plus and room for improvement and what people think should be improved varies by group and individual.

Second headline message is that the legitimacy belief towards ICANN generally correlate to your closeness to the regimen. The Staff has the highest legitimacy convictions, the Board comes after that, the Community comes after that, then if you go wider Internet Governance but people who are not involved in ICAAN it goes down and then you go to the wider world elites and then you go to the general public who have never heard of you and so they can't have legitimacy beliefs. That's the second message.

The third headline message is that more or less, there's a few exceptions but more or less, the legitimacy beliefs for ICANN are pretty steady across regions, across stakeholder groups, across social categories, gender, English speaking skill, age and so on. If you thought that there are big swings in support for ICANN< by region or by stakeholder group, the data doesn't show that. It's pretty steady. You could say it means there is no Achilles heel, there is no major weak point that could undo ICANN but at the same time you couldn't say there is a vanguard that is going to take ICANN forward to higher legitimacy.

The fourth thing to say, before we start into the details, this is descriptive data, you people are very technical, so you understand these sorts of things. We're showing you descriptive patterns, it's not for you to then immediately draw an explanatory inference from that. We haven't done a capsule analysis. We're telling you what the levels and patterns of legitimacy are but we're not telling you what causes that. That's what we're going to spend the next month doing and if

you can't put up with us after this talk, maybe we can come back in Cancun and tell you the explanatory stuff but that's still to come later.

Just to say what we mean by legitimacy, in academic language, it's the belief that a governor has rightful authority and exercises it appropriately. Or to put it more plainly, you think ICANN has the right to rule, has the right to make rules and people should obey them.

Legitimacy is important, you think so. We were kind of happy about this because if we were doing a study on legitimacy and you told us legitimacy wasn't important, we'd wonder what we were spending on it but in fact you say it's extremely important. The different lines are the Board, the Staff, the Community is yellow, the informed outsiders to ICANN is blue and the overall average is the blue on the left. What you're saying, extremely important vast majority, 80 percent of people say that it's extremely important for ICANN. Most of the rest, remaining 20 percent say that it's quite important.

RUSS MUNDY:

Is it just me or does it seem shocking that there is a Board member who thought it's only moderately important?

JAN AART SCHOLTE:

Well, you can register it. As you told us why you thought legitimacy was important, you said it was important because it helped ICANN to secure it's mandate, it helped ICANN to attract participation, to

achieve compliance, to hold off potential competitor institutions and other things. It's important.

Our evidence base, we talked to all 30 Board members between 2015 and 2018. We talked with 305 people from the various stakeholder groups in the community. We talked with 132 people from the Staff. By the way the figures that we'll show you later on are weighted for this, obviously there is an over representation of Staff for example there and actually Board as well, the figures later get weighted to reflect the actual proportion of participation in ICANN Meetings.

Then as for the outsiders, we talked with people not involved in ICANN but very involved internet governance and we talked with 860 general elites around the world, in six countries, in six world regions. We didn't do a public opinion survey. After many years of asking students, MA students whether they'd heard of ICANN and these are presumably very intelligent people and as each class at most has one person who has ever heard of ICANN, I'm presuming that people on the general street also don't know about ICANN. We didn't think it was worth doing a public opinion survey on ICANN's legitimacy.

Confidence in global governance institutions, this is looking at ICANN's confidence amongst world elites compared to other global governance institutions. You see ICANN there in the red line and the blue line is national government.

There are some happy conclusions that you could draw from this graph and there is some less, it's kind of half empty and half full. Let

me tell you the half full part first. Confidence in ICANN is generally at the similar or higher level compared with 10 multi-lateral inter-governmental institutions. It's also ahead of national government. The elites, those who knew about ICANN, believe in ICANN more than they believed in their nations state. ICANN, we asked about four new forms, private forms of global governance, multi stakeholder forms of global governance and ICANN comes out top amongst those four. That's the half full part. You can ICANN attracts quite a bit of legitimacy among world elites in that respect.

That average of 1.7 is close to the midway point on the scale of three. In fact, only the world health organization with 2.1, is attracting quite a lot of confidence, all the others are falling between quite a lot and not very much. It's not a resounding endorsement in that sense. But you could say, there's not really a resounding endorsement for any governance at any level, which maybe says something about our current time.

Another qualification is to say 49.7 percent of elites did not know of ICANN or didn't answer the question. After 20 years of ICANN's business, half of world elites don't even know that this place exists. You can't have legitimacy belief in something that you don't know.

UNKNOWN SPEAKER: Do we have ITU somewhere?

JAN AART SCHOLTE:

ITU, that's coming up. Then we did -- yes, we looked at confidence in various internet governance organizations. We actually asked people in the ICANN specific to say what they thought about ITU, IGF, ICANN, RIR's and IETF, and here you see that the IETF and the RIR's in general attract a higher average confidence in ICANN. The IGF and ITU and nation government attract a lower average confidence than ICANN.

Maybe, if you look at the overall pattern, you can say, the more a state is involved, the less the confidence, which is kind of interesting.

The dark blue lines at the far right of each of these groups, these are the informed outsiders. Here you should think these are people who are involved in internet governance but not involved in ICANN. You can see that their assessments of ICANN are notably lower. They still have the same pattern of the less state, the more confident and that might be a peculiar characteristic of people who get involved in internet governance or maybe it goes wider, we don't know.

Now we're going to ICANN much, much more specifically and then it's time for Hortense to take over.

HORTENSE JONGEN:

Thank you. Here we asked a question about how much confidence respondents have in the current workings of ICANN overall. We present a breakdown for the ICANN Board, ICANN Staff, ICANN Community and again, the informed outsiders, as well as the total of all of the insiders. As we interviewed a lot more people, relatively from

the Board and Staff, we also weighted these three groups, the community members have more weight in the total.

If we look at these results, again it depends on whether we look at this as glass being half empty or half full. On the one hand, on a positive interpretation is that when we look at the total insiders, the light blue, we can see that about half of respondents have high or even very high confidence in ICANN but when we look more specifically at the ICANN Community, we can see that more than half as very low, low or large majority moderate confidence in ICANN. The same goes for the informed outsiders, where the largest share only has moderate confidence in ICANN.

What we did next is we converted these answers to the survey question into numerical scores. If respondents indicated that they have very low confidence in ICANN, this was given a score of one and very high confidence in ICANN is given a score of five. Then we converted the mean averages for each group and what we can see is that ICANN Staff on average has the highest confidence in ICANN, they report a mean score of 4.11, which means between high and very high. This is followed by the ICANN Board who report high confidence in ICANN.

Then we have the ICANN stakeholder participants 3.45, this falls somewhere in between moderate and high confidence. Then again, we have the total of the insiders, the general elites, if converted that scale it's 3.27, again more towards the mid-point. Finally, the informed outsiders at an average of 3.18, this is more towards

moderate confidence again. What this shows is that the more deeply you are involved in ICANN regime, the more confidence you have in it.

This slide we present the breakdown per stakeholder group. We asked questions about confidence in ICANN overall as well as confidence in the ICANN Board, confidence in the multi stakeholder community and confidence in ICANN Staff. We found rather little variation according to stakeholder group, with a few exceptions. First of all, when it comes to confidence in ICANN overall, we can see academic has the highest confidence in ICANN, although this is based on a relatively small number of respondents, followed by the government group. Then we have the business other group who have a little bit less confidence in the ICANN Board and even less confidence in the multi stakeholder community compared to the other groups.

Finally, we have the government group who has the highest confidence in ICANN Staff, compared to the other groups, what is also interesting is even higher than in multi stakeholder community and the Board and in ICANN overall. As you can see the green bar is the technical community falls more or less in the middle when it comes to all stakeholder groups.

We found a little bit more variation when we look at the regional distribution of responses. Again, first we show you the average, which is 3.54 and then we can see that respondents from Russia and Central Asia, they report the lowest confidence in ICANN at an average of 3.05, more towards moderate. Although again, a relatively small number of respondents. Then respondents East, South and Southeast Asia, they

report the highest confidence in ICANN at 3.83. The difference is almost not 0.80 which is quite a lot.

Again, what is quite interesting about these results is that we cannot speak of a so called North South divide here because if anything, the respondents from Latin American and the Caribbean, Africa and Asia, they report higher confidence in ICANN compared to respondents from Europe and North America.

Finally, we compared social groups. First of all, we looked at gender and we didn't find difference between men and women in their confidence in ICANN overall. We also only found very little variation when it comes to different age groups. We also look at English language skills. Again, we found very little variation but interestingly native English speakers give lower average confidence compared to people with medium to no English skills. Actually, people with nonnative strong English skills, they report the highest confidence in ICANN.

Finally, we looked at staff reported race and ethnicity. We found the lowest confidence amongst people who identified as white and the highest confidence among Hispanics. Again, when it comes to these social groups, I would like to mention that these are just descriptive, later on we'll look into causal explanations for these.

Now, if you look more specifically at the responses provided by the technical community constituencies. We asked the question, in principle, so regardless of whether ICANN achieves the mater in

practice, how important do you find it, that ICANN achieves 15 specific aims. This again was on a one to five scale, where a score of five indicated that was extremely important and a score that is not important at all. Here find that transparency is commonly considered the most important for ICANN. This is about what ICANN should be doing, at an average of 4.86, which is a very high score, extremely important. Followed by promoting technical stability, taking decisions on the best available knowledge and expertise and finally accountability. The scores provided by the technical community constituencies are more or less the same as the other stakeholder groups.

When it comes to the four aims that are commonly consider least important for ICANN, promoting democracy in DNS management, promoting human rights in ICANN operations, promoting human rights in the DNS and promoting democratic values in wider society. These aims are commonly considered amongst members of the technical community from less importance than the average of all other stakeholder groups.

Then we moved to the question of, how ICANN is perceived to perform when it comes to achieving these aims. The technical community constituencies are generally most positive about ICANN's ability to promote technical stability, followed by technical security, to the extent to which gives all stakeholders the opportunities to participate in policy making and finally, that it takes decisions on the available knowledge and expertise.

When it comes to the four aspects that's perceived to do least well, it is promotion of human rights in the DNS, promotion of fair distribution of cost and benefits, DNI, promotion of democratic values in wider society and finally, taking decisions in a timely manner.

The question arises, to what extent is this problematic? When it comes to aims of promoting human rights, promoting democracy, as we can see on the previous slide, it's also not considered that important for ICANN to do, so the scores about the question, should ICANN be doing this, are also lower for these specific constituencies, other stakeholder groups might consider it more important.

It's interesting when it comes to taking decisions in a timely manner because the score of 2.51 indicates that ICANN is only able to do this to little extent or to a moderate extent, falling somewhere in between. Whereas when it comes to how important the technical community constituency finds this, the score is 4.34, which means it's between quite important and extremely important. Here we can see; expectations are not met.

As I mentioned earlier in this presentation, we have only covered descriptive, we've only shown levels and patterns of legitimacy toward ICANN and during the next steps we will also look into causal explanations to explain why certain groups hold certain perceptions to ICANN. Finally, drawing on these explanations, we could think about possible reforms in order to increase legitimacy perceptions, higher. If you are interested in this, we would be very happy to report on these issues at ICANN67 in Cancun. Thank you very much.

ANDREW MCCONACHIE: Is there any historical data you can compare this study to?

JAN AART SCHOLTE: That's the first question we get in the other presentations as well. I wish we did. No and it's a real problem because you have multi stakeholder global governance and it's been going for several decades, it would be great to be able to track what has happened over time, does it become more accepted as time goes on? This is the first study of its kind, so this has to be the baseline and then we have to come back and do it again.

ANDREW MCCONACHIE: You probably have historical data from the other institutions that aren't like in internet governance land but nothing for the institutions that are in internet governance land.

JAN AART SCHOLTE: Unfortunately, there is very little. The elite survey that we did there is unprecedented, so the general elite survey is unprecedented, it hasn't been done before. There is some public opinion data, in the world value survey but that's public opinion data and of course, as I said, ICANN just doesn't figure on the public conscious, we don't really get helped by that.

We have a little bit of longitudinal data on bodies like the United Nations and European Union but beyond that, we really don't have it. If you want to put this in a general context, why this is important is, global governing becomes more and more important in the way we live, so we need to understand what makes it tick. We're trying to catch up.

FRED BAKER:

I would image at least part of that is changes in the structure of ICANN, kind of in our community in general, the GAC didn't exist when ICANN was first formed, it was something that was added later because the other government didn't like the US running the show. There are other aspects of that but I would imagine part of that reflects changes that have happened in our community.

JAN AART SCHOLTE:

It would be wonderful to have data before the IANA transition to compare it with now but unfortunately, we don't.

RUSS MUNDY:

Are you going to make these slides available so we can at least where you are now with the study?

HORTENSE JONGEN:

Yes, we can send the complete presentation, we can send it by email, sure.

OZAN SAHIN: If you go to the ICANN66 meeting schedule for this session, you can find the slides there, I uploaded them.

JAN AART SCHOLTE: Ozan, we actually have a version that has a lot of our other commentary as well. If people wanted to have the full commentary, we can give you those slides. We can send you both.

PAUL HOFFMAN: Your last few slides, I think what you were saying because it wasn't in the headline, was the technical communities' views of ICANN, the importance and how we were following through, is that correct? Did you do the same type of analysis for other organizations, such as some of the ones you'd mentioned earlier? The ones I'd be interested in is ITF and ITU or is this really about ICANN? Is your greater study really just about ICANN?

HORTENSE JONGEN: Yes, thank you. Indeed, we only asked these questions about ICANN specifically. We have questions about confidence in other internet governance institutions that we presented but we didn't as them specifically to rate their performance or procedures.

PAUL HOFFMAN: Not that I get to set your agenda but often times people will say, ICANN is or is not like so and so else, and it would be wonderful if you were looking for other topics; some of those other organizations are just as ripe for these kind of considerations as ICANN are.

Certainly the ITU, which does have a longer history and there have been some other academics who have looked at the IETF, not on the legitimacy scale as much, although some of their studies really did come into the legitimacy question and certainly some of your questions at the end, which I'm not at all surprised that ICANN got -- in fact I'm surprised we got even a 2.5 from the technical community on promoting human rights because the technical community care about promoting human rights by in large but the IETF might have some. If you're looking for other targets, I would say some of the other technical communities, most of the kinds of things I saw here, if you were giving the same presentation at ITU or the IETF, I would expect you would have similar high results for at least responses.

JAN AART SCHOLTE: Thanks for that, I'll note it. We have put in a research proposal to do something similar about RIR's but it's very hard to get research money, so I don't promise.

WES HARDAKER: First off, I thank you for your work and I certainly sympathize with getting research money, I fight that all the time being an academic. The work was fantastic and fascinating, I appreciate it. There was a lot

of cases where the results were actually quite similar and there was a lot of cases where they were quite different, did you guys calculate error margins and how close some of those -- did you calculate error margins and how close some of those actually mattered?

HORTENSE JONGEN: Indeed, I don't present a standard deviation, I haven't looked at whether the responses are very scattered within groups. For some papers we're doing comparison of mean test and we'll also after that run more sophisticated regression analysis but this only at this stage, we've only coffered descriptive statistics.

FRED BAKER: I was interested in the divisions of different countries and their views of things, North Americans, South Americans and so on and so forth, did you try clustering the countries and these are similar and these similar and different groups of similarity, as opposed to geographic proximity?

JAN AART SCHOLTE: We didn't ask people for their country, we asked them for their region. We can infer sometimes that maybe Canadian opinions on certain issues would differ from US American opinions but we have no way of actually set that out. It was part of the anonymity, to try and ensure the anonymity of people, if we started to ask them their specific country, then in many cases we would have only one or two people

and that would go against their anonymity and we wanted to make sure that we had confidentiality and anonymity in the responses here.

I think there's an interesting issue to think about in terms of the region, seeing those much higher scores for Asia and Africa and Latin America, relative to Europe and North America, it's somewhat counterintuitive in some senses, if you look -- what you might assume to be patterns of influence and benefit from the ICANN regime but those inequalities actually if they exist and we have other work where we ask people whether they perceive inequalities, but the result and the confidence scores is actually the opposite. There is a question in my mind to say, why are those scores from Africa and Asia and Latin America considerably higher than the European and North American scores?

FRED BAKER:

One particular cluster that I would expect to find in the data, if you had countries in it, would be the BRIC countries, Brazil, Russia, India, China, Indonesia, South Africa; they form a block that tends to think that what we do here should be transferred to the UN or to national governments, and actually in a discussion where Brad and I made a presentation yesterday to the GAC, to the Government Advisory Committee, we had a specific comment from China, and China is very much one of those that wants this transferred to national entities. I would expect those countries to be a common opinion. That's where I was going with clustering the countries.

JAN AART SCHOLTE:

In the general elite survey, the target countries included Brazil, Russia and South Africa, not India and China in this case, so I can't say something about those three. They're scores for ICANN were actually the same as or higher than US and Germany or a little bit lower. It makes you ask, is it about a particular part of -- for example the Russian government, where as Russian elites in general are actually much more relaxed about this. Maybe.

There is also the question, when you talk with people from the GAC, they are of course a particular part of the government and they may not actually correspond with what the foreign ministry and cabinet and the president want. Again, the GAC numbers, amongst the stakeholder groups, the GAC average confidence for ICANN was the next to highest, which again in some respects, might be counterintuitive because you think the business constituencies and the technical constituencies might feel that this is more their place, if you had certain assumptions. This data doesn't suggest that at all.

HORTENSE JONGEN:

If I could add to your question about clustering of different regions. We do have data on East Asia, South Asia and Southeast Asia also separately. Our motivation to merge this data when we present it here is because the responses didn't differ much. The one reason for example not to put together the Middle East and North Africa and Sub-Saharan Africa is because the response is different. We look at similar

responses and the clustering of specific regions but indeed, as Jan Aart just mentioned, we couldn't look at specific countries and how we wanted to cluster them.

FRED BAKER:

We're pretty much out of time. Does anybody else have a question or comments they want to bring up before we break? Seeing none. Thank you.

JAN AART SCHOLTE:

Thanks very much, thank you.

BRAD VERD:

We're back here at 1:30 to talk metrics.

[END OF TRANSCRIPTION]