KIM CARLSON: Hi, everyone, and welcome to today's NCAP Discussion Group call on January 27th at 19:00 UTC. In the interest of time there will be no roll call. Attendance will be taken based on those on Zoom. Kathy and I will update the wiki with the names of the participants as quickly as possible. We have apologies from Barry Leiba and Matt Larson. I believe Rod Rasmussen is also an apology.

All calls are recorded and transcribed and that recording transcript will be posted on the public wiki. Again, as a reminder, to avoid background noise and echoing while others are speaking, please mute your phones and microphones. And with that I'd like to turn the call back over you, Matt. Thanks.

MATTHEW THOMAS: Thanks, Kim, for that. Appreciate it. Welcome everyone to our weekly NCAP Discussion Group call. I hope everyone's doing well and still staying healthy. On today's agenda, we plan on continuing our case studies. Today we're going to hopefully cover .corp as well as .home. And then, if Jim is able to join us later in the meeting, he will hopefully give us an update on the Study 2 project going forward.

So, with that I'll kick off item number 2 in our agenda, Updates to SOIs. If anyone has an update to their SOI that they'd like to declare at this time, please raise your hand. Not seeing any hands.

So with that why don't we just go ahead and get into the two studies. If we can start on the .corp one, please, that would be appreciated. And

maybe I'll copy and paste that link into the chatroom for everyone. So, this will be a similar analysis like we've seen on .mail and .internal but for this string we're going to take a look at .corp. So, hopefully most of these graphs look familiar to you. There will be a few new ones and a few ones that are removed, that I'll explain later going forward. If we can go ahead to the next slide, please.

So, this is a longitudinal view of .corp from A and J root servers looking at total traffic volume on a per day basis. The left side is split up by root: A-root, J-root, and old J-root, as well as splitting it up by IPv4 and IPv6. But the graph on the right is the aggregate of those three root letters. You will notice that in March, or the first quarter of 2020, .corp traffic significantly went up. In fact it almost tripled overnight in the middle of March. And as we all know, that's really when the pandemic started here this last year.

So, one theory might be is that all of these devices that were traditionally in their corporate networks were now being used in residential or non-corporate networks. And this .corp string that should have been resolved within the confines of that DNS environment prior, is now going out into the global DNS. We'll take a look at that and maybe confirm that or support that with some additional data going forward. On the next slides, please.

So, again this is a similar plot looking at traffic volume. But again, at this time we're looking at the specific QTYPES. This is very similar to what we saw on .internal as well as .mail. There is nothing I would say particularly interesting in the QTYPES as with something like .arpa. However again, here we're seeing a larger percentage than I would say

normal delegated TLDs per SRV records inside of .corp. And this is DNS query type that's commonly associated with zero configuration or DNS service discovery type events. If we can go ahead to the next slide, please.

This, again, is looking at a longitudinal plot, but this time looking at the number of unique IP addresses requesting .corp. And here again, in the early March or early 2020 timeframe, we see that the number of unique daily IPs hitting A and J goes up roughly 40-50%. So, suddenly .corp is coming from a much larger, much wider-dispersed amount of sources out there on the Internet. Again, timing this with the COVID event, maybe this is likely, again, those transient devices leaving their corporate networks and being used in more residential places. And if we take a look at the next slide, please.

This is just looking at a global distribution. Where is .corp actually coming from? You can see almost half of the traffic is coming from the US. There's some traffic coming out of France and India, but then it quickly dives into the long tail around the rest of the countries around the world.

If we go to the next slide, please, and we start to look at the networks that are sending this traffic, it starts to get a little bit more insightful here around us. So, if you remember compared to .mail, where it roughly only had 900 and some ASNs, .corp is now spread out over almost 5,000 ASNs. If you look at the cumulative distribution function on the left there, to get to the 95th percentile of all traffic, you would roughly have to "remediate" or interact with 250 different ASs to understand and fix this .corp query leak out to the root.

But the graph on the right is looking at the top ASNs sending .corp traffic. While I don't have their names listed down there, if you look at their AS numbers and look them up on the Internet, almost all of those are residential ISPs, either the United States or Europe, specifically France. So again, this is a little bit more evidence that this traffic now is suddenly coming out of those types of ASNs, before where it was maybe in the corporate environment. If we can go forward to the next slide, please.

And I think this is maybe even more interesting part to understand what's actually going on in .corp. The graph on the left is a histogram plot of what percent of the names contain how many labels. So, roughly 10% of the names coming in for .corp only contain .corp, but the vast majority of them, I'd say 60-65% combined, have three or four labels. So, these are much more "rich" QNAMES that we're seeing come into the roots that probably give us a better sense of what's going on with the actual source of those queries.

And if we take a look on the two lists on the right side here, the first list in the middle is looking at the top 30 SLDs ranked by the total percentage of traffic caused for that. If you go through the various lists of names in there, it starts to become apparent that many of these are indeed major corporations like .airbus or .zuerich. And it would be a logical assumption that they've just elected to use this .corp string within their corporate network.

However, you start to go down the list and you also see things like .internal, which in my mind, you would have either elected .internal or .corp and maybe some of this is also still the by-product of suffix search

list appendage and .corp is getting appended to things like .internal there. So, you have company X.internal.corp. ending up being on there.

And some of this seems to be representative. You can see that on the third label list on the right—we're at number 29—you see this IHS.internal.corp, right? But the rest of the third labels definitely give you the sense that these strings like .zuerich, where they have an additional label in front of that, where it's like AMER. It's anchoring that within a particular region like AMEA or the UK or so forth and so on. So, this is a unique aspect of .corp in terms of QNAME structures themselves that we can observe and understand from this. If we can go forward to the next slide and the following, please.

Here we're going to take a quick look at the data sensitivity analysis. And again, just as a reminder, this is trying to get … At the core of what we're trying to look at is if we look at one root, or if we look at two roots, or if we look at three roots, how much of the overall picture do we have? And so here, I had split the SLD overlap on a Venn diagram on the left down by A-root and J-root for the last day of December 2020. While there is a descent amount of overlap on the SLDs, clearly each one of the roots has its own unique capture point.

And this is reflected again longitudinally on the graph on the right where the overlap or the flattening of that curve—the blue and the green lines—doesn't really happen. They seem to still be going at a linear growth rate.

If we go to the next slide, please, and we look at possibly why that is the case, the graph on the left is a cumulative distribution plot of how many

times a specific SLD was queried under .corp in the entire month of December. And here, you can see roughly 80% of these strings or SLDs are only being queried one time. So, this is probably something indicative of some kind of random label or some kind of jitter in there. Maybe this is Chromium.

That makes this measurement of understanding. Do you have a complete snapshot to assess the risk of the actual overall collection of SLDs is accurate and complete? The graph on the right is then just another growth curve except I eliminate all of the SLDs that have only been queried once for the month. And here you can see that the green and the blue lines at least start to flatten down a little bit.

And likewise, if we go to the next slide, we're going to repeat that analysis on the amount of ASNs or the different various recursive resolvers or network operators out there that we see at A and J root and combined.

The Venn diagram on the left, again, is showing that A and J have a fairly significant overlap, but J-root has 1600 different ASNs that it sees but A-root never saw. So again, that catchment of a particular root and its contribution to the data analysis is somewhat telling. And I think maybe this is one of those questions that we're going to want to further study when we do the data sensitivity analysis, by either asking additional roots to help fill out this Venn diagram or we conduct that kind of exercise on something like the DITL data. Unfortunately, things like the DITL data won't allow us to compute graphs on the right since we're typically limited to two days. However, it at least would allow us to figure out the overall overlap between all of the different roots.

So again, that graph on the right then is just looking at the longitudinal growth of distinct IPs querying for .corp, here at least you get to start seeing some of the flattening of the curve over the period of the month.

But again, if we continue to the next slide, please—thank you—many of these IPs, as you can see on the left cumulative distribution plot, send roughly 10 queries over the course of the month to make up 30% of it. So, the graph on the right, again, is trying to figure out if that curve flattens anymore if we subset that data. I don't think there's anything exceptionally insightful from that to be honest though. At that point, I think that concludes the overview of .corp. Did anyone have any questions or thoughts on .corp at this time? Not seeing any hands and that's totally fine.

So, why don't we switch over to .home because that will give us at least something to compare to. At the end of this presentation, I actually created a few more slides that combine .corp, .home, and .mail all together on one slide so we can "compare" the curves and talk about the differences that we see between these three different strings.

One of the things I'm hoping that these case studies are doing is it showcases that each string is not the same. Each string has its own "unique footprint" and its own "unique fingerprint." And so, I'm hoping that by the end of this and we've talked about those side-by-side graphs that we'll be able to maybe help understand some of that in terms of the questions that we think should be asked when conducting risk assessments and how we fold that into providing guidance for the questions, ultimately, that we need to answer. So, if we can go to the next slide, please, Kim. Thanks.

Okay. Again, standard longitudinal graph looking at .home this time from A and J. Again, in early 2020, in the March timeframe, we saw that query volume almost went up 2.5X from what it previously was. But then, suddenly, at the end of November into December of 2020, there's a pretty rapid decline. And I have some slides going forward but this decline is due to the change in chromium and the change in the way it sends out probe requests on the Android platform to detect the Internet redirector. So, a pretty noticeable decrease in .home there on that. The next slide, please.

Again, just taking the same look at the data longitudinally but this time looking at QTYPES. I don't think there's anything extremely eye-catching here. Again, it's mainly A and quad A, but you will note that SRV records make up significantly less than .corp here. So, maybe that is an indicator of how the .home string is actually being used. I think as we go through this .home string it's going to become apparent that .home is actually probably caused from suffix search list appendage, where .corp was probably intentionally anchored by those various corporations. Next slide, please.

And so, this again is taking a look at the number of unique IP addresses for .home over time. And here you can see again early 2020 the amount of IP addresses for this string significantly increases. While it does decrease a little bit near the end of 2020 with the Chromium change, it still remains relatively high compared to its earlier level in 2020, likely because everyone's still at home working from home or remotely. Next slide, please.

So, here we're taking a look at the traffic distribution for .home, and .home seems to be not nearly as US-centric as .corp. The US does make up for 30% of it; the rest of the world does seem to be more heavily utilizing .home. Specifically, countries in Europe, and China, and Asia, seem to be a little bit more heavy users in them. But overall, again, the traffic is globally dispersed throughout. Next slide, please.

And I apologize for some the axes being so difficult to read. The CDF on the left, the x-axis actually tops out around 9500 or a little bit past 9,000. So, there were over 9,000 ASNs that A and J were receiving traffic for .home. This is significantly more than anything that was coming out of .corp and .mail, so .home is very widespread. And as you can see, to get to the 95th percentile, you're going to probably have to talk to upwards of 500 or so ASNs. The top ASNs are plotted on the right, I didn't label them again, but I can tell you that all of those are major Internet service providers in the US, as well as Europe and throughout the world. Next slide, please.

And again, I think this is one of the more interesting things in .home compared to .corp and .mail and this is the actual label analysis. So again, on the left we're looking at the number of labels present in the actual strings.

What I found interesting was that almost 75% of all the queries were coming out containing only two labels, which seems a little bit strange to me. I'm not quite sure why, especially when you take a look at the list in the middle at the second label domains. One thing to me jumps out is that many of those second labels are actually delegated top-level

domains. They're either country codes or legacy or new gTLD domains in there that all seem to have this .home randomly appended to it.

So, I think this is very different than what you saw under .corp, which those strings fundamentally look like they were purposely anchored onto as corporation or this entity. Whereas here, it seems that .home is just getting appended to these various standard TLDs that the world would interact with in the DNS ecosystem.

Then, when you look at the column all the way to the right where it's looking at what's underneath of those second label domains, you start to see some of the most common domains out there on the Internet—things like Google APIs or Google.com, so forth and so on. Oh, yes. Sorry. Warren, please. I didn't see your hand. Go ahead.

WARREN KUMARI: No worries. I figured I'd say it in the chat as well. And I'm muted so you didn't hear me typing.

MATTHEW THOMAS: That's awesome. I need a drink of water anyways.

WARREN KUMARI: So, this is just a random supposition, but what it kind of feels like for the reason why there are second levels of delegated names, is it kind of feels like there is some sort of CPE which is integrated software which is trying to do QNAME minimization and doing it badly. Like there was a dnsdist release which tried to imprint QNAME minimization and it looks

like maybe it is QNAME minimizing, then doing the search list, and then trying to do the rest of the label length. I wonder if that could be it? I'm not sure if what I said made any sense?

MATTHEW THOMAS: Yeah. That's a good theory. So, dnsdist or something like Dnsmasq could be having—

WARREN KUMARI: Sorry, masq is what I was meaning.

MATTHEW THOMAS: Oh, okay.

WARREN KUMARI: Yeah, there was a Dnsmasq version where somebody tried to hack in QNAME minimization and did it really badly, it got released. And it sort of feels like the QNAME minimization part might be being applied, then the search list, and that ending badly. What would be useful would be to be able to hunt down somebody who's got a piece of CPE, who's doing these lookups, and beat them and see if we can run it in a lab.

MATTHEW THOMAS: I think that's an actual great idea. And I think your theory on this is probably correct, because if you look in the middle column, like number seven, the Fios router, that's definitely a CPE device from a major ISP. Not saying that they used the Dnsmasq implementation that you're

referring to, but I would say that that's a good clue that this is an avenue well-worth digging into.

WARREN KUMARI: Yeah. The Hitron or Hitron I believe is also an existing cable modem gateway thing, which is a fairly awful one. I guess I shouldn't say that while being recorded, but a fairly awful one.

MATTHEW THOMAS: We're a small group. There we go. Yes. That's another great example of that. And I think number 21 as well, right? UniFi, isn't that Ubiquiti?

WARREN KUMARI: Yep. UniFi is the Ubiquiti built-in-home thing. Looking through here, I think that there's a lot of things where CPE are just annoying. And as you saw from the distribution from the ASNs, if you have a look, many of the ASNs were home network vendors like Sky in the UK and Fios.

MATTHEW THOMAS: Yep. Exactly. And I see Jothan is mentioning ISATAP and WPAD in the chatroom. Yes, those are definitely Microsoft-based. And on an upcoming slide, I'm actually going to change from just looking at the second and these third label domains, but to looking at what is on the first label of the QNAME. So, I think that is something interesting, maybe, that we can take a look at as well. Kim, if we can go to the next slide, please.

So, again this is just looking at the properties of those second-level domains coming in. Now, the CDF on the left again shows that 85% of these are only, again, receiving one query. And you might be thinking where's that graph on the right that shows the longitudinal growth over the course of the month from A and J, and the Venn diagram, and everything? Well, one day of QNAMES coming in at A and J for .home was over 322 million, and that actually broke down to 277 million unique SLDs. And my poor little laptop did not seem to like to load that into memory. So, I didn't really go down the avenue of trying to do all of that and do MapReduce. But I think the graph below explains what we're probably experiencing under there.

And so, this is taking a look at those SLDs based off of the length of characters of that string. And so, you see that kind of flat bar between 7 and 15 characters long. There is the spike at 9, I believe. Or is that 10? I can't tell. But again, these are likely all just Chromium queries. So, again the vast majority of this traffic coming in under here is something related to that and probably not worth the cycles for me to pull my hair out or make more of a [grade] to make that a longitudinal plot. However, if we feel that there is need, I will try to do that. But if we can go to the next slide, please.

And this is jumping to what I was talking about before. So, the table on the left are the most popular first labels contained under .home. So, this is the very first label in the QNAME that's specified. And I thought this was interesting because it shows you really what some of the things that are being used in here are for. So, number two, the underscore, we know that is associated with some kind of QNAME minimizing recursive resolvers out there. WPAD, the web proxy auto discovery, we know is

set up of for zero configuration. That is a concern from looking at name collisions.

Then you have other standard labels that are associated with common, I would say, Internet protocols and/or just various different things you would interact with like www. or API or LB for load balancer. You have the ISATAP in there. And then again all of these. As you go down, you can start to see some of the Apple caches. And there is a, I think, convention for DNS service discovery things to prefix that first label with an underscore.

And so, to that end, I made the column on the right which is the most common first label strings that start off with an underscore. And so, I thought this was interesting to see what service discovery protocols were running under .home. A lot of them are under apple. But then you have other things like Bradford Agent and some of the PCOIP bootstrap. And doing a little google foo; you can find all of those strings out there. And they are associated with other various types of either CTE devices or other kinds of network devices that utilize DNS service discovery to bootstrap. So, I thought this was useful because it indicates what is actually being used under .home and it seems like it's a pretty wide varying amount of stuff there. If we can go to the next slide, please.

Again, this is just trying to do the data sensitivity analysis, but like I mentioned before, the 277 million SLDs kind of killed the SLD growth curve. So, I only performed this on a network IP and ASN basis. Again, that Venn diagram on the left is looking at the overlap between A and J root. Again there's a fairly significant overlap. But again, J-root for whatever reason, still sees 2500 more ASNs than A-root did. And then

looking on the right you have that growth curve. They don't seem to be really falling out at all over time.

Again, that Venn diagram on the left is looking at the overlap between A and J roots. There's a fairly significant overlap, but again, J root for whatever reason still sees 2500 more ASNs than A root did. And looking at on the right, you had that growth curve. They don't seem to be really falling out at all over time.

If we continue to the next graph, please, this is where I wanted to take the three case studies, specifically corp, home and mail since we know that is a specific question that the Board has tasked us to look at specifically and put some of these measurements together on one graph so we can kind of compare apples to apples.

The first one, again, is just looking at query volume. You can see here on the left that is a query volume. The one on the right is looking at the number of unique IPs.

In terms of query volume on the left, .home is obviously several magnitudes larger than mail, maybe a magnitude or more higher than corp. That being said, we know that query volume always is an indicator of risk, but it also kind of gives you a sense of which way the wind is blowing, maybe.

The graph on the right is looking at the number of district IPs. The nice thing here is that you see that same kind of uplift across all three strings during the early portion of 2020, which we're kind of associating with the coronavirus and remote transient devices.

But one last comment on the graph on the left, I think it's just really awesome to see that dip in .home from Chromium. I think that's such an awesome, amazing move. That's very encouraging, I think. Next slide, please.

So this is looking at the number of ASNs for each one of those strings and looking at their distribution profiles. And I think this is kind of important, more important than query volume, because this talks about, in my mind, "the footprint" of this particular string. Is this string super disperse and there's traffic coming from it all over different places, or is it relatively confined to a set of particular networks or ASNs?

And I think that this is important because it helps inform your decision making and your thought process about risk assessment and remediation possibilities.

So if you looked at .mail on the bottom, they have less than, I think it was 900 different ASes, and we said that—I forget off the top of my head, but it was like if you did the top ten or something like that, it was going to remediate more than 50% of the traffic, whereas you can compare that to something like .home where it is spread out over a much larger network or a network operator basis and to get that kind of remediation or control of that string, you're going to have to interact with a whole lot of different operators, unless you can find clues within the QNAMEs or other signals that we're seeing in the data to help identify the underlying cause and trying to do a remediation via that mechanism. Next slide, please.

That is really the kind of point that I was trying to make here on looking at those second-level domains and the additional label information that you see in each one of these strings, is that if you look at for instance .home, we saw that most of those second labels are what appear to be delegated TLDs or things associated with CTE devices, and you compare and contrast that to .corp where those are clearly different corporations or businesses that have anchored themselves under .corp versus something in .mail on the right, you look at that and I see much more of those names in there, just as standard kind of protocol indicators, NS1, NS2, SMTP, things of that nature.

So you have three very different observation points there to consider. And I think looking at the [inaudible] like this is it gives you an insight into maybe what it's being used for and also help you inform what's the possibility of being able to identify the underlying cause and remediate it. If this is truly caused by CPE devices and you can figure that out, then how long does it take to work with those CPE devices and push that out and remediate that? And I think that's another thought or concern here on how we conduct a risk assessment and provide guidance. We have to weigh these costs between the traffic, the properties of the DNS and so forth.

I guess with that, that was where I was kind of leaning towards, is that what I'm calling the footprint, given the footprint or how big a particular string is based off of network diversity or ASNs or whatever, and its "fingerprint," AKA the types of queries or particular types of QNAMEs that it's sending, is that risk manageable? What options are there to remediate it? And if it is yes to either one of those, how can that be done, in what kind of timeframe?

And I think that's where I was hoping that these three case studies would at least show the differences between them and also how, maybe, we need to take the underlying lessons from these and apply it more broadly.

With that, I've been talking for a while. Does anyone have any other questions or comments at this point? Not seeing any hands. I do hope this has been useful. I think the plan going forward will be to continue doing a couple more case studies. I know we identified, I think, two or three more strings in the study two, so my plan is to hopefully have those strings ready and presented for next week, and then we can continue this analysis so we have a larger bucket of fish to compare, and then we'll just keep going with our study. I did see that Jim Galvin has been able to join, so if he's available, I'd like to turn it back to him and give him an opportunity to update the group on study two and how it's progressing.

JAMES GALVIN: Sure. Thanks, Matt. So the status of study two is we had one—the revised study two proposal, we had one comment from an SSAC member that raised questions about conflict of interest and the way we're going to continue addressing them. Given that we've kind of changed the dynamics of the relationships within the project, there were some questions about the relationship of that to what was originally specified in the original proposal.

Now, just to be clear for folks here, in this project, in this discussion group, we have adopted the ordinary ICANN working model of

statement of interest and we just expect that those in the project have a balanced view of those interests, and that's why everyone here has filled that out, it's published on the ICANN website, as is ordinary, and we also have—it's either seven or nine, I forget now but there's a few extra questions that we ask specific to NCAP just to more directly expose more potential interests that folks might have with respect to this project.

So yeah, if you weren't remembering that you had a separate statement of interest from the ordinary ICANN SOI on the community Wiki, you might want to take a look at that again and make sure that you're up to date as far as that's concerned.

But we're going to publish that soon. We'll have that detail together. There's one comment that's going to be SSAC's normal publication process allows for indivudals who have a point of view separate from the consensus of the group. They can elect to have that included as part of the publication. So we're just going through a little bit of internal process to let that person express their particular concerns that they have about conflict of interest and statements of interest in this project. That will be added to the project, revised project proposal that you've all seen. It's going to remain where it is.

Sorry, I'm just pausing for a moment here because I'm thinking to myself we actually added some text to the proposal that we have to send around to this group. Actually, I thought Steve Sheng sent it here. It already went out to the group.

So we added a little bit of text explaining what we—redrafting section 3.2 from the original proposal. We added a little bit of explanatory text in response to this comment, but the comment is still a concern. So they're going to add a dissent to the proposal that's going to be published. I'm sorry, I probably just jumbled all of that up for you. Let me pause for a moment, take a step back and try to say again what's going on here. But I see Warren has a hand up. Why don't I let Warren speak for a minute, give me a moment to collect my thoughts? Let me try to say that again much more clearly so people can get it. Go ahead, Warren.

WARREN KUMARI: Thanks, Jim. I will happily admit I got sidetracked while you were saying some of the initial part. So you probably already covered it. But you said there's an additional set of SOIs which we all submitted, and we should double check them. Could you provide a pointer to those? I think you might have mentioned it but I lost track where it was. Just so we can make sure that they are up to date. Or I misunderstood you. One of the two.

JAMES GALVIN: No, you're absolutely right, Warren. Thank you. It's actually on the ICANN community Wiki, the NCAP project page. There's some links in there. I'll put a link right to it into the chat room here.

WARREN KUMARI: That would be great, because navigating the community Wiki is challenging.

JAMES GALVIN: Yes, it's its own joy. I know sometimes it can be kind of interesting. So let me step back. There was one comment in SSAC, there was some discussion and some concern about potential changes that are present with respect to conflict of interest, given that our revised proposal kind of changed the relationship among the people doing the work, moving to a project sponsor, contractors, the kinds of contractors, the discussion group doing more of the work.

So what we did is we did add some text to the proposal, which you should have seen, to respond to those comments, and in response to the text which exists in section 3.2 in the original proposal. So it simply reframes what's there to reflect the fact that we're operating now much more under ordinary ICANN working group kinds of processes. But there's still some concern about some of the details, and so what we're waiting for is some additional—the individual SSAC processes allow for someone to include an extra statement of a personal nature and they can simply state what their concern is, and that will be added to this document and it'll be there. But we didn't really want to make any particular change to the document and go through yet another review process and such. We were more inclined to just explain the changes that we believe happened as a result of the way we've restructured the project and then otherwise remain operating under the same SOI, COI mode that we were operating under.

The BTC meeting, the Board technical committee meets tomorrow, and we're going to give a quick presentation to them, a heads up about what's coming. In particular, we'll be showing them a draft timeline and also a picture of the financial side of it, what's drafted in there in terms of what's there, which, the financial side of it, this group has not seen, but that's by design. Only the admin committee works and develops with ICANN the financial side of the project proposal.

So we'll be presenting that to the Board tomorrow, and then as soon as we have a finalized document, we'll send it here, the discussion group will get a copy of the final document, and then we'll be sending that to the BTC and then we'll be on their timeline. Hopefully—we expect that they'll want to approve the project going forward and the support that we're looking for, and then that will kick off ICANN processes to look for the roles that we're looking for from contractors, and we'll get all of that support.

So I'm sorry, Matt, I didn't intend to be quite so longwinded about that. I got a little confused on the first round. But if there's any questions, happy to take those. I hope that was clear enough to people. Thanks.

MATTHEW THOMAS: Thanks, Jim, for that. Really appreciate it. With that, I think we're at item number six, Any Other Business. Jim, please go ahead.

JAMES GALVIN: Yeah, and I'll jump in real quick just to say thanks to Kathy for putting the link in the chat room that Warren had asked for. That's where you

can go to check your own SOI to see what's there. So, thank you, Kathy. And Warren, just let us know if that doesn't answer your question. Thanks.

MATTHEW THOMAS: Thanks, Jim. Anyone else have Any Other Business at this moment? Otherwise, I think we can wrap up today and you can all get 12 minutes of your day back. We will be back here next Wednesday to look at the next set of strings. Okay, well, let's call it a meeting then. Thank you, everyone. Have a good week.

[END OF TRANSCRIPTION]