
KIMBERLY CARLSON:

Thank you and welcome to today's NCAP discussion group call on March 25th at 19:00 UTC. In the interest of time, there'll be no roll call. Attendance will be taken based on who's on Zoom. Kathy and I will update the Wiki with the names of the participants as quickly as possible. Looks like Warren's on, so we just have an apology from Richard Wilhelm who will just be late.

As a reminder, calls are recorded and transcribed, and recoding transcripts will be posted on the public Wiki. Also, to avoid background noise while others are speaking, please mute your phones and microphones. With that, I'll turn the call back over to you, Jim.

JAMES GALVIN:

Thanks very much, Kim, and thanks to my co-chairs, Matt Thomas and Patrick Fallstrom who we have with us. Welcome back, everyone. I hope your travels to and from Cancun weren't too stressful for you. At least that joke still works in this group, or sort of joke, I suppose.

Anyway, yes, we are now going to start having weekly meetings and we are going to try and dig in as much as we can to continue moving forward with the work that we have in front of us.

We have a quick agenda here. Oh, I'm already forgetting the formality. Since this is an open ICANN group, we have a little bit of extra formality of checking updates to SOIs. Has anyone made an update to their SOI that they just want to call out to the group here?

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.

Not seeing any hands or hearing any voice. And we don't have any new members at this time, which is maybe a little disappointing, but hopefully as we start to get in here a little bit, I've always felt like the subsequent procedures group was taking away some interest in this group as people were focused on wrapping that up, and we'll just see and kind of move on for ourselves. Next up here is to check on the status of public comments.

I don't see Matt Larson on the call, but as far as I know, there have been no public comments on what was put out, and it's not closed yet, it's open for another week. But to the extent that the work product one that Karen has produced so far in all of this is really just a bit [biography.] I'm not really expecting any public comments, unless somebody for some reason has some sort of objection to what Karen wrote about our document, but none of us have had any issues with it. Maybe they have something that got left out or something. That would be nice to get, I suppose. But I expect that to close with very little fanfare and then Karen will wrap that up and Matt will do his part to ship that off to the board, and we will have that for ourselves too.

It will then be up to us to take on those various references and use them where appropriate as we continue through the rest of our work. And Jim Prendergast is saying here in the chat room that he did check and there were no public comments in the call for it, so that's a good thing.

So with that, just a reminder about the board questions. That is actually a link to our—actually, you called that board questions, but that doesn't look like the link—oh, it is. That's the link to the spreadsheet that we've

been using up to this point to go through our discussion. We had copied out each of the board questions that they had asked of us in the resolution, and there are ten items there, they're numbered one to nine with the first one in the list just being just the separate resolution which was a reference to .corp, .home and .mail. And thanks to Kim there who put that up for folks to see what we're talking about here.

You'll notice that we have Column E, which is our discussion points for ourselves that we were creating about each of those items, and then there's actually a Column G. again, for each of those questions from the board, we had gone through and we've had some discussion, and we have identified what we believe are some gap questions, meaning some additional questions that we'd like to ask that at least so far, we believe this new data that should be examined as part of looking at the questions in Column A, and we were trying to capture some points over there in Column G.

So that's just a reminder in that particular agenda item. The next agenda item is really where I want to go in this particular group here, which is a document called work product one gap, and it really is a collection of each of the things in Column G. That really is all that's there. The boldface is simply the item from Column A, the reference to the board comments, and then inside of that is the bullet point notes that I had taken while we were having the discussions at the time and had put there. so that was just to capture our thoughts in that moment.

And Kim helpfully put the pointer to that doc in there. That one should be accessible to all the panelists. I don't know offhand if others can see it right away. I don't mind if it's in view mode for really anyone, but

panelists should be able to edit. And if you make any changes—and I do welcome folks to just type stuff into the document as usual, but please change your mode to suggesting just so that we can track what's changed, and of course, we get a look at who did it so you don't get to make anonymous changes to the document. Sorry about that.

Okay, so what is it that we're after, and what are we doing? Let me take a moment to just put this in context, and then we'll see if anybody has any questions or comments. We have study two and three are really about the process by which we're going to go about answering the board questions, and it's up to us to create the substance that's going to help us provide the recommendations to the board so that they will be able to speak to the issues that they gave us to talk about and asked us to write about and research and comment on.

So the work product one at the moment, part of what's happening is that the board made a decision to create work product one and we went through an RFP process, ICANN did that. We have Karen here with us and she did work product one.

There is a second part of her work product in which she has to offer a recommendation to the board as to whether or not it's appropriate to continue and fund studies two and ultimately study three, but right now, the next step in the process will be to fund study two to collect some input and some facts based on the bibliography and provide a recommendation to the board on whether to continue to fund this group.

Our opportunity here—and that’s what this work product one gap analysis is intended to represent—we’ve already collected some notes. We can certainly add notes to other items if folks have stuff they want to add. Like I said, I just pulled out what we had pulled out before.

This is our opportunity to consider what additional things, what additional questions, what additional analysis we believe we would like to do that have not been done before. Are there questions or issues that we’d like to explore with data that we can hopefully get our hands on and thus motivate the need for being able to fund some analysis work that will help us down the road do the actual data driven analysis so that we have evidence to work with as we do our analysis to answer the issues that the board has put in front of us.

I hope that that’s clear. I feel like I’ve kind of said it a few times over, but we’ve had kind of a gap here since we’ve met last, so I wanted to take an opportunity to walk through that a little bit, bring everyone back up to speed, make sure that we’re all on the same page here, and also talk a little bit more specifically about exactly how we’re going to proceed.

The goal here is just to look again at these notes and consider these questions. Is there more that we can say here? What we ultimately have to do is convert this into some kind of text recommendation, something that we can give to Karen that she will then craft into a response to the board. And hopefully, she’ll be working with Matt Larson and OCTO, office of the CTO at ICANN to create what is going to be the recommendations to the board, but hopefully, our work here will be of significant influence there. So the goal is for us to have this discussion

and provide this to her so that she can fold this into her analysis and her job and her role at producing the work product.

I'm going to pause there for a moment. Does anyone have any questions or want to add any comments? Jaap, welcome. I'm sorry that you've missed all of my lovely introduction. You'll have to go back and catch up on the recording later. Steve Crocker, your hand is up. Go ahead, please.

STEVE CROCKER:

Yeah, looking at the gap analysis, a thought occurred to me. Apologies if this has been discussed before. I was thinking about a general rule that ICANN might add to the requirements for new delegations, and even for existing delegations of the form that a TLD registry should be obliged to treat—I'll call them stray queries in a responsible fashion.

So that's kind of a notional thing, and then we can unpack that a little bit. Whether or not it's a heavily used one that has an awful lot of traffic, like .corp, .home or .mail, or .belkin or something, or whether it's just an ordinary new TLD that gets some stray traffic, I can imagine putting in a requirement that says if you're an operator registry and you're going to get either innocently or maliciously bad traffic, you should have a way of dealing with it which does not make the problem worse.

So if it's a query to, say, star.newtld, you don't give an answer to that, or we have some specification what to do with that. And if it's to a nonexistent second-level domain, here is the way you should be dealing with it.

JAMES GALVIN:

Thank you for that, Steve. Let me ask a clarifying question here. I tried to capture the point that you were making in one sentence up there as a bullet point in the document. Feel free to edit that or expand on that in any way that you like. But that feels like it's more of a place that we want to get to in our recommendations later. Is this something that we need now? How would we use that now in terms of studying data in order to motivate study two with the board?

STEVE CROCKER:

Yes, I think this is where I want to get to, and so I apologize if it's out of order, but I thought I would share it while it was on my mind. It does provide a little bit of guidance with respect to data to collect in that it would be interesting to know what the nature of the queries are. We have a lot of data on that in .corp, .home and .mail, what the second level domains are and so forth, to see if there's anything that could be said about traffic that shouldn't be there.

Should there be—I guess I'm just going to keep riffing on this. I could imagine that we'd learn something by looking at what the experience has been. We'd also learn something about, is it possible to detect patterns? And again, then jumping ahead to what the recommendations might be, should there be obligation to look at the pattern and begin to take some proactive or seek help on getting proactive remedies?

I'm just going to break off mid-sentence there, because either that stimulates some thought, or it doesn't.

JAMES GALVIN:

Thank you for that. I like the comment that Anne is making in the chat room and certainly would welcome any response from folks on what you said there. Anne had made a comment in the chat room about maybe data that's needed is to take a look at data after the 90 days of controlled interruption.

And that's true, and although I think one of the things that we certainly want to look at which is included in the bibliography is what we know about what happened during controlled interruption. So yes, if anything was noticed during controlled interruption, I agree, we certainly should try to take a look at what did happen. That's a sensible comment as far as I'm concerned.

Of course, now, we're talking about stuff which was eight years ago, and maybe some of that data is not generally available anymore. That would have been an interesting question to ask a long time ago. But let's at least make note of it for right now. I'll copy this comment over and put it in this document here for the moment. It's certainly a question worth asking as to whether any of that data—if any root server operators in particular happened to keep that data around for that long.

Okay, I captured these two extra points at the moment to the top. I don't really have a specific place to stick those two comments. So, any other comments or questions from anyone about what we're trying to do here and what our process is? Otherwise, this might turn into a fairly short call if we're just going to review our notes here. We have to figure out how to formulate these comments that we have already made to

ourselves, beginning with label two, and turn them into something that might be useful for Karen to include. Or maybe, Karen, I don't know how much you've studied this, but as you get into this, if you have questions for us that you would like for the group to consider, please do speak up and let us know. We're happy to have discussions about whatever would be helpful to you.

And Matt Larson is commenting about the only source for old root server data is the OARC's DITL, which is only 48 hours every year. All true. Well, it might still be worth just taking a note and looking at that. I'll certainly make that note here at least so we have it. Maybe when the time comes, we can see if we can take a look and see if there's anything there. It may not be that there's much there, but it's probably worth a look nonetheless, just in case something pops out looking at NX domain data in TLDs that have been delegated over the last eight years. Warren, go ahead, please.

WARREN KUMARI:

It feels like for many of these sorts of things, we should also be thinking of next time we do this, we should do a better job of collecting this sort of thing. Like as part of the initial name collision stuff, there was a discussion that maybe instead of just counting them or something, registries should be required to keep a list of these and report better stats. So maybe if we do this again, there should be a "registries have to keep a list of them in case there's a third or fourth round or an ongoing round" or who knows, and we won't be in the position of going, "Wouldn't it be nice if we'd collected this last time?"

JAMES GALVIN:

I'm going to make a note about that up here at the top too. That's a good suggestion for continued analysis as new rounds proceed. We certainly should have that opportunity to do what we can, and we should think about what the board might want to ...

This is a way to evolve the recommendations that we do. I'm thinking about this a little bit and just imagining that whatever set of recommendations we ultimately can provide to the board, they're going to want to manage those recommendations too. They might evolve or need to evolve, and perhaps being able to look at recorded data over time for decisions that have been made would be a good thing.

Also, keep in mind that one of the questions that we have to answer is about the board used the phrase, "Collision string," and, "A list of collision strings," meaning those things which are identified to be in a category that should not be delegated. And is it possible that things on that list might ultimately be reevaluated and be allowed to be delegated because something may have changed, whatever that is?

Seems to me that this bullet point of items here about what data can we continue to collect about strings, this falls into that category. Especially if something goes on to the collision string list, then yes, we're going to want to be able to look at that and continue to watch it, and maybe that'll be the basis for moving off the collision string list in the future. Something to think about.

Okay, any other questions or comments? And I apologize, it sounded like somebody was speaking out earlier. Does somebody want to jump in again here? Matt, go ahead, please.

MATT LARSON:

Yeah, I just wanted to piggyback off Matt's comment about DNS OARC's DITL data. Verisign's retained A and J NXD traffic for numerous years now, so we have full fidelity data, I think going back at least four, maybe five years that we'd be able to make use of for any analysis that we wanted to look at. That's obviously not going to cover all of the new gTLD delegations, but it might give us a subset to look at.

JAMES GALVIN:

Okay. That's good. Thank you. And I apologize, just trying to make some notes here in real time as we go along. And Ram asked a question in the chat room, don't we need to figure out the criteria for a collision status to be removed?

And the answer to that is most definitely yes. Part of what we have to do here in examining the list of issues that the board gave to us is to consider the question of how to determine what is a collision string, what kinds of criteria can we give to the board to review so that it can make the decision that it needs to make about whether something is there or not. And then the other half of that is presumably we'll be able to provide some guidance about when that data is sufficient to be considered not on the collision string list. So all of that is something to be seen. So yes, thank you.

And Ram, your hand is up. Go ahead, please.

RAM MOHAN:

Thanks, Jim. Just following through on that, I guess the corollary point I'd like to make is while the criteria for a string to be called collision and therefore placed on some sort of a protected status, while that may be defined or definable and set up, it feels to me like the criteria to remove a TLD from that status may not just be a mirror image. In other words, if there are six steps that mean that a string is placed in a collision category, the absence of those six data points, we have to think about whether that is sufficient evidence to have a string come off that list, number one.

And then number two, is this something where once a string has been placed on a collision list and then comes off a collision list, do those strings need some level of monitoring of data? Is that required? And is there potential or possibility that if such a string comes off a collision list and is then delegated to some party, that there is a risk that that string may at some point in time go back onto a collision list?

In other words, once you come off a collision list, is it irrevocable and therefore allows ICANN to delegate, or do we need to think about some cooling off period or waiting period? How do we determine what it is? It seems like the removal of the status might require quite a bit more work than placing on that status.

JAMES GALVIN: Thanks very much for that, Ram. I'm making a couple of bullet point notes here in the document. Please feel free to expand or update this if I don't capture this quite right here, but I hope I'm capturing the points that you're making there in the document. Thank you for that. And I see Warren is agreeing, many plus ones.

RAM MOHAN: Jim, sorry, I think on the first sub-bullet where you say separate work item, I think the key point is not that it's a separate work item but that it's not a mirror image. In other words, if there are six steps to get on a list, it may be that the absence of those six metrics may not be sufficient. So the point I was making was it's simply not a reflective thing. One side is reflected on the other.

JAMES GALVIN: Okay, I switched that up to your use of the phrase "mirror image." Okay, thank you. Yes, and as Warren says in the chat room, absence of evidence is not evidence of absence. Okay, all very true too. We're dealing with a "We don't know what we don't know" kind of thing here. So all good.

All right, thanks very much for these comments. Glad to see that people are engaged and adding some thoughts here. I think that, just to keep this discussion focused a little bit, I made all of those points, I tried to capture them up at the top, but so far, I view all of these points—and folks, please correct me if I'm wrong or you want to shape this differently so that it gets us to a different place—I view these points here as part of our analysis discussion that we ultimately have to have.

We're going to have to go through kind of a brainstorming analysis discussion and looking at various circumstances to get to the recommendations that we ultimately need to produce here.

And the focus that I have right now is really to think about what we wanted from study two when this full project was originally laid out was about what data can we look at, and what new data can we examine, what old data can we reexamine that has not been considered before that will help us in characterizing this collision string problem space and the issues highlighted by the board?

We need to identify that those things are present, and the questions that we might want to ask of that data. That's what we need to get and prepare. It probably doesn't have to be really expensive, but we do have to focus in on a few things so that Karen has that to include, and then essentially, we're trying to motivate a need, if it's there. obviously, if it's not there, it's not there, but ideally motivate a need to do some continued analysis like what was done back in 2011 and 2012 in front of the last round, which got us controlled interruption. Do we need to do that again? And if we do, what are we going to get out of it that wasn't there before?

One comment that I would make is something that has always been on the table in my point of view in this work is the desire to look at the public resolver data if possible. if it would be possible to ask similar questions that were asked of root server data back in 2011 and 2012, I'd like to ask similar questions of resolver data now. I think that there's a new set of data there.

I think it would be interesting to compare root server data to resolver data just to see if there's been any real change, and it'd be interesting to ask a few of these larger public resolvers if they're seeing behavior similar to what we've seen by root server behavior earlier. We need to consider if this kind of scenario exists today, and also to think about whether or not it's possible to work out some kind of relationship with these resolvers so that perhaps the presence of collision data—I'll call it for the moment—is something that could be available for making this decision about delegation or not.

I'm just going to let that sit for a minute. I see that Ram put a question in the chat room here, and he asked, do we have evidence that controlled interruption was a successful mitigation model for name collision?

And we only have a handful of controlled interruption reports, and we only have a handful of reports from controlled interruption before. Part of the study two analysis was really to review those handful of reports that we do have and actually try to dig into them a little more and really do some research with the parties that were affected and see what we can learn from that, do a real root cause analysis and understand the circumstances that created that.

So you're right then, we could fill out the question of what metrics do or don't work and whether or not controlled interruption was successful or not. We have very limited data, quite honestly, because we only know about a handful of reports. So by any statistical measure, that's just anecdotal. But nonetheless, we should at least take what we have and look at it in detail and then maybe we can further drive the future of

controlled interruption. If that turns out to be the recommendation that we're going to go forward with, we can further influence what is or is not included in that. And yeah, thank you, Warren, absence of evidence is not evidence of absence again. Keep us all a little lighthearted here.

But yeah, I guess as Ram is saying in the chat room, just to stay on focus here and stay on point. We don't really know if there's any evidence or not there. We at least have a few things to examine. So part of study two would be to take those particular elements that we do know about and dig in and do whatever it takes to find out as much as we can about what happened there so that we can write about that and study it and see, and then perhaps learn about some additional metrics that we might use going forward to make this a standing system, if you will, for continued evaluation and management.

This feels like the Jim monolog. Good, Ram, let's hear another voice besides mine.

RAM MOHAN:

Thanks, Jim. One concern is that we have a precedent here about identifying strings and getting all strings to go through controlled interruption and also placing some strings on collusion string status.

If you look at having every string go through controlled interruption, in the last round there were several new TLD players who complained and said it was unnecessary, there was no problem, etc. And I think that I don't know that there was a constructive way to respond to a critique like that. That's really what I think they ought to be looking at.

If you look at the document that the JAS, etc. put out, there were some assumptions made about what controlled interruption will achieve and how to determine whether it is achieving that. So there were some assumptions that allowed ICANN board to make a determination to do controlled interruption and to have controlled interruption required for all strings.

I don't know that there's anything substantive that has been done since that—I forget, 2010, 2011—paper about this topic that might allow the board or the community to have better visibility into, is this a good mitigation mechanism or not? So that's really what I'm trying to get at. And to do that, we have to, I think, go back to the original controlled interruption paper and see what assumptions it made and what metrics it intended to cover, and then see if we can gather that data. We ought to have plenty of evidence between then and now.

JAMES GALVIN:

Thank you, Ram. The one thing that I want to point out is it was study three, the next study after this, which was intended to do a deep dive into mitigation strategies. I'm sure that we can do some of this work a bit in parallel, and we should absolutely do that so that we don't serialize everything here and push the end of this too far out.

But the model that we had when we had laid out the full project was that we would have this study one and in study two, we would do some analysis of the data that we had, and we would also consider potential mitigation strategies in addition to controlled interruption. We would begin to brainstorm some ideas as we learned about what the data

showed us, what we might be able to create in terms of how to manage some of that data or whatever we learned from it, and then the goal of study three was to do a deep dive into various mitigation strategies, including actually creating a bit of a lab, if possible, where we could test some of the strategies and see.

Now, part of our analysis has to include how the board might actually evaluate proposed mitigation strategies in the future. Proposed new strings might come with their own ideas for how to manage collisions that may or may not be present or may or may not come forward once their string is announced. So we need to give some due consideration to that. That's part of our own tabletop analysis exercise we're going to have to do in all of this.

So I tried to capture your point about controlled interruption. Yes, you're talking about data. This is actually kind of important. In the context of our explicit work item right now, although this is detailed work for the future, one of the things that's important here is to be able to get at that data. This is new data—just like the resolver data—which has become available to us. It's limited data, but there's definitely data there and we should take the opportunity to do a deep dive into that data and learn what we can from it, pull out of it everything that we can so that we can have that as input to our analysis about mitigation strategies in general, and of course, in general what collision strings look for.

So I want to frame your comment there in those two parts: what we can do now and what we can do in the future. Let me give you a chance to

comment on that, on whether you want to agree with that or whether you want to suggest something different.

RAM MOHAN: Yeah, I agree with that, and that's kind of the path I was suggesting or thinking of as well, Jim, which is collect or at least understand what is collectible now and see what it takes to start collecting it. Study three can look at mitigation strategies, etc., and that's a parallel track, but what data already exists that ought to start being put together? I think that's something that can be usefully done relatively soon.

JAMES GALVIN: The study was done by JAS before, right?

RAM MOHAN: Yeah, JAS did it before.

JAMES GALVIN: Okay. That's a very useful data point for our purposes right now and work that we're doing, so thank you for that. Matt Thomas, you have your hand up. Go ahead, please.

MATT THOMAS: Yeah, just off of Ram's note there on collecting data around the efficacy of controlled interruption, it just came to my mind that I believe numerous different pieces of software, specifically some of the browsers, have hardcoded in the logic to handle the controlled

interruption IP addresses. So it might be useful to see if they have any instrumentation statistics around that that they could share with us to help guide some of the efficacy measurements on controlled interruption and how often they're actually seeing it.

JAMES GALVIN: Thanks for that.

MATT THOMAS: I think Firefox is one.

JAMES GALVIN: Okay. Thanks for that. Warren, go ahead, please.

WARREN KUMARI: I believe that Chrome was doing some sort of detection on that and popped up a, I think, fairly useless error message. But it was not actually collecting any stats on that as far as I know. I did look into it and there was no summarized or aggregated counter on the number of times it happened. It was as purely local, no feedback, thing as far as I know. I think it was viewed that it would be creepy if it reported that sort of thing, even just an aggregate count. As far as I remember.

JAMES GALVIN: Okay. That's okay, but certainly, if you have the opportunity, you can double check that. If that changes, let us know. It'd be interesting to see

if there's anything there and what we can do with that. So thanks very much for that, Warren.

Wow, this has been good. We haven't even gotten to the questions yet. [Going down to data is there.] Maybe just scroll down a bit there, Kim, so people can see the bottom of what I'm writing there, the last couple of bullet points, and folks can keep up with what's going on here.

So yeah, I guess I'll just make a larger note here at the top about access to resolver data. Been saying that quite a few times, but might as well make a note of it here right here at the top. Let's say some words again here about what I said before. Just a point there to keep us on track so I can keep track of what I had said before. We can fill that out more completely later.

Okay, that's at least two data points, two chunks of data that we know were not examined before, because of course, it wasn't available before. That might be interesting to look at this time, and then compare it to the root server data that was gathered before, whatever statistics was done before. And as long as we're going to go that far, we probably ought to ask root server data even this time around. We should be able to compare to today's statistics, gather up the same kind of graphs and statistics, metrics that we had grabbed before in 2012, and just see if things are essentially roughly the same and then compare that to what we get out of resolvers, controlled interruption, and of course, the historical evidence that we have too.

Okay, great discussion. Thank you so much. I'm so glad that people had given this some thought and have some things to say. I'm looking at the

time here, being a little bit conscious about the time. We might just take a few minutes here and do a quick look through some of these questions and remind ourselves of some of this data. If you can pull number two there up to the top of the screen, Kim.

I think one of the things that's interesting here that jumps out at me is realizing that there is some additional data available now that wasn't available then, or at least some additional considerations, the notion of NSEC caching and QNAME minimization, and the local root discussions. These are all things that are evidence that the infrastructure in the broader Internet and its behavior has changed since 2012. That only makes sense. It's been seven, eight years. Of course these things are going to evolve. So it really is appropriate to consider these new technological behaviors about the infrastructure in the context of collision strings and see if that has moved what the collision string graphs look like, what NXDOMAIN graphs look like and where that change has occurred.

Maybe it's all the same and it's just moved to resolvers as opposed to root server operators, or there's some relationship there that's worth taking note of. Matt, you have your hand up. Go ahead, please.

MATT THOMAS:

Maybe to dive into number two a little bit more, speaking about the evolution of the DNS ecosystem since the last round, I think one of the things that's going on right now that's interesting in how it relies on a negative answer is Mozilla's implementation of DOH using a canary domain to influence the application to stop using the stub resolver over

into the application space. That's all based off of a negative answer from that canary domain.

And we've been observing at the root DNS suffix search list being applied to that canary domain and essentially "resolving" in unintended spaces up at the root under other nondelegated TLDs. So I think that's kind of an interesting data point for us to look into where the intended use of the DNS has transpired from a lookup into an application logic switch, but because of the complexities of DNS service discovery or suffix search lists, we have induced this unintended resolving of the name in an unintended space or a collision there, possibly.

JAMES GALVIN:

I want to capture the phrase that you used there, Matt. I just put it down here, my second bullet, "The DNS has moved from just being a lookup tool to an application context switch?" I don't think that's quite what you said.

MATT THOMAS:

It's kind of evolved from a traditional lookup, from a name to an IP address, to now being used in application to decide how resolver behavior is changing, that it stops using the stub and starts using the application DNS.

JAMES GALVIN:

Yeah, you used the word "context" and you used the phrased there, "to an application." It turned into context switch in my head. But what did

you call it? I forget what word you used, and I guess you don't remember either off hand

MATT THOMAS: [inaudible]. Sorry.

JAMES GALVIN: I'll just keep it as context switch right now in quotes. At least I understand what you're getting at. That'll remind me here for right now.

[BRAD VERD:] It'll be in the transcript.

JAMES GALVIN: Ah, yes, there's always the transcript later, come back to it there. Thanks for that. But you're right, actually, maybe that's the larger key point here. You added DOH here which is a good one, certainly has become quite a significant data point—DOH and DOT— in security discussions. But the larger, more general point that the DNS infrastructure itself, the Internet infrastructure really has evolved and the DNS in particular as an application, as evidenced by these few things that we have here, has really changed.

And it really would be worth doing the analysis that was done years ago again, because each of these things has had some effect on DNS behavior. No one's really looked at what that effect really is. I appreciate that some folks here—yourself and probably Warren at

Google and some others I'm sure—have done some internal analysis and look at these things, but in the context of collision strings, this is the new thing. And I think it's important to see how that might have affected—or not. It's still useful to know that there's been no real change, but documenting that and observing that would be a good thing.

Okay. Continuing on for a couple more minutes, just to see here, the question about harm, I don't think there's anything there that I wanted to bring up here. We had some helpful thoughts there. And I'm just going to scroll through these kind of quickly unless someone wants to jump up and say something here. I'm reading through them myself as just an immediate reply, just to see immediately if something jumps out at me as a way to capture a meta point here that we need to make in this influencing of study two.

And item five here gets us to the discussions about controlled interruption and other mitigation actions. We talked about that up above. We talked specifically about controlled interruption, but the future of mitigation is really the study three activity. There really is a question of whether there might be other mitigation strategies besides controlled interruption. It was the choice at hand eight years ago. We have to allow for the possibility that others might come up in the future and give some thought how to address that.

And then we have seven and eight where we want to get to the question of how you would be added to the name collision string, how you could be removed from it, and we did capture that up above again already for ourselves.

And then of course, item three, more colloquially, what this is really getting at—just so that we’re all clear on what’s happening here, this is a nice sort of unbiased technical description of the issue at hand, but I always like to characterize item three as this is about trying to check out gaming of the system. How do we deal with people who may try to game the system in some way? And that’s really the thing that we want to give some thought to.

One can never be perfect about these things, I suppose, but it is important to give the topic some thought and identify whatever we can. And whatever we can say about it that would be helpful to the future of evaluating collision strings and their presence or absence or movement on and off the list of collision strings.

Ok, that brings us to the bottom of this thing. This has all been very helpful from my point of view. I think the action here at this point is to try and turn this into a narrative that tries to answer the question whether or not to proceed with study two. And I guess our co-chairs here, Patrik and Matt, myself, we get to volunteer to make a first pass at that. I’d really appreciate—if anyone wants to volunteer to provide some text and help us in crafting this, that’d be great, and then I’d like to say that we’ll certainly have something for next week’s discussion, and then we’ll take the opportunity to see what this group has to say about what we write and then we’ll hand that off to Karen and give her a chance to see how she’s going to make use of that and let her draft what she needs so that she can then come back and share a draft of her work product just for our review and we’ll see where that takes us from there.

I think Matt Larson, we don't want to drag this out too long, of course, but I think that puts us on a timeframe of you and Karen being able to produce something approximately by the end of April. I don't see any reason why it should extend past that. And maybe just a few weeks from now, all of that should be together, and so the rest of the process with public comment and whatever that you need to do with that can happen. Do you want to comment on that at all, Matt? You don't have to. So we'll just leave that there for the moment unless you want to speak up.

MATT LARSON.

Sure. The plan all along has been to be done by June 30th. The public comment got started a couple of weeks late, so I think it's possible we'll bite into July, but we're not materially behind the schedule that we had planned.

JAMES GALVIN:

And Anne's asked me a question in the chat room: does the recommendation to proceed study two need to come formally from both of SSAC and OCTO, or just SSAC, or what? And actually, no, that's part of the project that OCTO and Matt Larson are managing. So it's just part of the work product that they're producing and that they're going to deliver to the board. So that's sort of the formality of what's going on there. And it doesn't even formally come from us, Anne, in this group. Obviously, we get to provide our input to all of that and the expectation is that our influence will align with what Matt and Karen and OCTO is going to come to. So we're just going with that plan for the moment.

We're not going to try to turn this into anything other than what it needs to be, and just sort of keep things moving forward.

I think with that, I'm going to ask for Any Other Business, and in particular, make sure that my co-chairs have an opportunity to jump in here if they want to add anything in closing.

MATT LARSON: All good here.

JAMES GALVIN: Okay. We will meet next week. Hopefully, we'll have some semblance of a narrative. I'll try and turn some of this into some kind of narrative that we can use here to move forward with, and so we'll have some discussion about that next week, and then we'll consider how we're going to move forward with our more detailed analysis going forward after that.

So thanks, everyone. We're adjourned.

KIMBERLY CARLSON: Thanks all. Bye.

[END OF TRANSCRIPTION]