BRIAN AITCHISON: …of safeguards that we have discussed testing. I intend for the session to be quite interactive, because we'll want to build out this matrix here to determine which safeguards we want to test and how we want to test them. So I'll walk you through the table and show you how it works and how I've set it up. Then we can go through and start some discussion as to which safeguards we want to focus on, and how we want to test them.

You'll see here that we have several groups of safeguards. We have from the DNS Abuse Report. We have – I'm not sure how many – I believe six, seven stemming from Spec 11 and GAC advice. And we have several pertaining to rights protection safeguards, which you can see here. It's about ten of them. There's also one on name collision. I know that's something that has been a focus, so we'll want to discuss how to study that as well.

But first, I want to get your feedback on these broad categories of safeguards and where you think our focus should be, if there's any sets that we should eliminate. So, for example, I think we're all fairly on board with looking at the DNS Abuse Report safeguards. Is that safe to say? I think we can just move right on from that.

Now, getting a little, perhaps, more difficult, is GAC advice and Spec 11-type safeguards. Are there any issues with including that within the scope of what we study? Any objections to not including them?

Go ahead, Laureen.

LAUREEN KAPIN: Sorry, you said it one way, and then you said it the opposite way and I got confused. So are you proposing to include the GAC safeguards and the Spec 11 issues, or to exclude them?

BRIAN AITCHISON: I'm asking, yeah. I am proposing to include them. I think they're an important aspect, and they're not really addressed elsewhere, from what I can tell.

LAUREEN KAPIN: Right.

BRIAN AITCHISON: In a rigorous way. So I think we're all on board with keeping this in our matrix. And you'll see, why I'm asking this is because of this third set here, these rights protection safeguards.

Right, go ahead, Carlton.

CARLTON SAMUELS: Do they link?

BRIAN AITCHISON: They link up. The issue, in my mind, is that they're already being extensively addressed in other areas, including a Rights Protection PDP. So there might be some sort of redundant work being done there. Do

we have any thoughts on that? And this is all tentative right now. This is for you, essentially.

Carlos?

CARLOS RAUL GUTIERREZ: Yes, I have a problem, looking at it in these three groups. And we have a basic assumption in the Charter about consumers, and we [divide] to consumers and applicants and final users. That's our boundary. And I would like to go through the list and see if it fits into these two groups, or in other reviews, I agree with you. But not sticking to the strict three groups that you have. If in the end we have groups and they look the same, excellent. But before I can say yes, I would like to go through them, analyzing from the point of view of our perspective. And our perspective is not Spec 11, Spec 13, but applicants, application process, up to the signing the agreement, and then use. Are those safeguards more oriented to protect the user during the use, takedowns, and so on? Look more at the cycle than at the way you have said it. And as I said, it might be exactly the same, but I cannot say those, yes, and those, yes, and the other ones, I don't know. That's my comment.

BRIAN AITCHISON: Okay, great. That's useful.

UNIDENTIFIED FEMALE: So if I could jump in, because you raise a good question, Carlos. But because I anticipate that that sort of inquiry, at least for the half hour or so we have now, will not be able to take place in real time.

CARLTON SAMUELS: We could [inaudible].

BRIAN AITCHISON: Sure, absolutely.

UNIDENTIFIED FEMALE: Yeah, that we go through that inquiry, just not during this time right now.

BRIAN AITCHISON: Okay. No, that's actually very helpful, Carlos, and that's the kind of input I'm looking for, is how to categorize this, what's useful to you, in terms of what you're looking for.

So, great, let's jump into what we have here.

KAREN LENTZ: Brian, sorry?

BRIAN AITCHISON: Yeah?

KAREN LENTZ: Can I just add a comment on the rights protection?

BRIAN AITCHISON: Sure.

KAREN LENTZ: So this is something that we haven't talked a lot about in this group, in terms of the trademark protections that were built into the program. And if you look back at the scope as described in the Affirmation of Commitments, that third item is effectively the safeguards that were built into the program. So I just want to make sure that, in terms of planning out the work, trademark protections was a very key issue and key set of safeguards that were built into the program. So I think it's key to make sure that that category is addressed in the review. Thanks.

BRIAN AITCHISON: Any other comments in that vein?

Okay, so let's get into the nitty-gritty of it. I'll show you just how I've categorized all this. So let's take the first one, vetting of registry operators. We have a safeguard. We have a high-level method, a way to address it. It really only lends itself to a qualitative inquiry. It doesn't really generate date to conduct a statistical analysis. So I've labeled it a qualitative study. Under source and method, you can see that I've emphasized certain areas where a vendor might be more useful and certain areas where ICANN might be able to provide research on it. And then some notes for discussion.

And all this, if you look at the far-right column, are decision points based on what I've called a "bang for buck" index, where we will get the most bang for buck, in terms of money spent, resources, and whatnot. It's a

bit of an eyeball measure at this point, but you can see point 2 here essentially breaks it down into the meaningfulness of the results we could achieve, the amount of research legwork, the sample size we have, and the methodological expertise that would be required to conduct a study of each safeguard. And that sort of filters into this, what I've called BFB index, where we can have a high or low bang for buck hiring a vendor or not. And you can see, I've put my two cents here, just to sort of get a discussion started. And we can see what you think.

So that's how it works. So jumping into this vetting of registry operators, ICANN, we do have some reporting from Pricewaterhousecoopers, who conducted the background screenings. It's not very comprehensive. It's high level. A lot of it is sensitive, so we can't delve too much into it. But we can possibly provide enough on our end that it would not necessarily necessitate a vendor.

Alternatively, we could take on a vendor to do something like a "perception of effectiveness" survey, questionnaire, focus group, set of interviews, something like that. That's what I've labeled it, since we are testing effectiveness. So I wonder, do you have any thoughts on method or studies, ways to study this safeguard, beyond what I've already suggested here? And if we're okay with what I've proposed, then that's fine. We could take that as – yeah, absolutely, please.

UNIDENTIFIED MALE: So you're referring to the cohort that is now a part of the program, not the ones who've been dismissed or who've… Do we have any

information about any registry operator [inaudible] who didn't make it because they had a bad flag from PwC?

BRIAN AITCHISON: [inaudible] there's been no terminations of any registry agreement. And as far as I know, no one was excluded from registering based on the finding of the [background]. There's not [inaudible].

UNIDENTIFIED MALE: Okay, so there's no real…

BRIAN AITCHISON: As you can see, there's a low bang for buck in hiring a vendor.

UNIDENTIFIED MALE: Yeah, nothing. It wouldn't be useful.

BRIAN AITCHISON: Okay.

DREW BAGLEY: For that one, I'm wondering though what we could come up with to determine whether or not the safeguard itself could be improved nonetheless. Because if the bar was set so low so that everyone would pass such a background check anyway, then perhaps it wasn't going to be an effective safeguard anyway.

BRIAN AITCHISON: Yeah, that also gets into the issue of [remembering the turn] effect. It's difficult to [inaudible] it might not just ever find.

DREW BAGLEY: Right.

UNIDENTIFIED FEMALE: I'm going to remind everyone to identify themselves and use their microphone.

BRIAN AITCHISON: Karen, I think I saw your hand go up.

KAREN LENTZ: Thanks, Brian. So just to remind everyone on the data in relation to the vetting of registry operators and the background screening is fairly limited, because we're working with personally identifiable information and had certain terms around how that data was going to be used. And in terms of the reporting, because the individual case reports are confidential, you don't know whether, for example, some of the withdrawals were as a result of this process or at what stage what iterations needed to occur before it was clear whether an applicant met the threshold or not. So we're kind of limited in that sense, and so it's hard to come up with some sort of numbers to study.

But to Drew's point, I think you also can look at – assuming that you'd want there to be some form of background screening to continue, I think there's other areas that could maybe be examined. For example, I think if you talked to applicants, you'd probably get some feedback about the amount of information that was collected. Did it make sense? Were they asked to update the names and personal information of their board of directors? Was it the right amount of individuals? Was it too many? Was it too few? That kind of exercise might get you some more input as to the effectiveness of the background screening. Thanks.

BRIAN AITCHISON: Carlton?

CARLTON SAMUELS: We had a situation where just about everything that they offered – it wasn't a whole hell of a lot – was on a confidential basis. And we are now saying that we are going to ask them to tell us. Well, did we ask the right questions to tell you it was confidential or not? I'm not sure I understand the impact of that.

LAUREEN KAPIN: Karen, jump in if I misunderstood you. I didn't take Karen's suggestion as bearing on what you just said, Carlton. I understood it as saying that we have some limited data regarding this safeguard. Because we have limited data, and because it may be hard to really draw firm conclusions about it, there isn't just one conclusion to reach. There could be a range. Karen is suggesting a way to gather more information would be

to speak to the applicants and delve further, at least into the issues of, "Do you think these were the right questions to be asking you?" Not saying we should ask more questions or get into more confidential information, but just ask the people who actually had to provide this information, "Hey, you provided this subset of information? Do you think this was the right information? Or do you think if you want to get at a credentialing or screening process, really the key information is something else?" That is what I thought I heard Karen say.

Karen, jump in and let me know if I misunderstood.

KAREN LENTZ: No, that's correct. I think you can explore qualitatively people's impressions of the way it was constructed.

BRIAN AITCHISON: And this is great. This is the intention behind this discussion. We might want to move the focus that I've put onto a vendor if we think that a survey of applicants would be useful. It sounds like it could be.

UNIDENTIFIED MALE: I was just going to say, I think this safeguard analysis will also be informed by that overarching study we keep referring to on DNS abuse. Because if we hypothetically found that one specific TLD had an overwhelming amount of abuse over others, and it was agnostic to the registrar selling the domain name, we might be able to ask ourselves if there is something particular about that registry operator. And so that would definitely not tell us anything definitive in and of itself, but that

would also help inform our look at the safeguard and what recommendations we might make about the effectiveness of the safeguard or research that needed to be done specifically on the safeguard, proposed research.

BRIAN AITCHISON: I've captured some of those notes. I'll kind of want to move through these relatively quickly. Is there anything else that we want to say about this one? I'll send out a summary, as well.

Okay, let's move on to DNSSEC deployment. This one seems like a pretty important one. It seems like we have quite a bit of quantitative data on it. It's something that could be theoretically tied to some abuse rates. So there's a lot we can play with there. And also, I think enlisting the help of a vendor would be quite useful with statistical and data-gathering expertise. Do we have any thoughts on that?

UNIDENTIFIED MALE: I didn't get it, if that was more internal because you have the whole data, or if you need a vendor.

BRIAN AITCHISON: We would need a vendor, because it would likely require pretty significant statistical expertise, I would imagine.

CARLTON SAMUELS: So we are looking at DNSSEC deployment on the gTLDs. Okay. And is that not a part of the reporting requirement now?

KAREN LENTZ: On DNSSEC, there's a few areas that you can look at. All of the New gTLD registries are required in their agreement to implement DNSSEC and to post their – it's called a DPS statement, I believe, some information that describes their DNSSEC in the TLD. But part of what DNSSEC does is it creates this chain of signatures. And going beyond the top level, if you look at how many second-level domains are signed and how many other aspects of the registration chain DNSSEC occurs in, that's kind of outside of our scope of knowledge, or at least direct knowledge, although there are statistics and reports on that out there. I think there's a little bit of a difference, in terms of what ICANN has when you're looking at the registry level versus the rest of the ecosystems, Calvin.

CALVIN BROWNE: It would be possible to use CZDS to get the number of signed zones in each zone. It wouldn't be too difficult to do.

UNIDENTIFIED MALE: I was just going to say, if you go to the website rick.eng.br/dnssecstat, that's the site referred to by the Internet Society on DNSSEC siding. And the site's updated daily. So I think whatever our bigger study is on DNS abuse, we could then get that vendor to correlate their data with these statistics, or we could do this pretty easily ourselves. I don't think we

would need a separate vendor for this category to answer these statistics.

BRIAN AITCHISON: Right. I think that's a good point. That site is something we referred to in the DNS Abuse Report too. I think it'll end up being pretty useful.

And also, just to be clear, the way I've envisioned this is sort of a one, full service vendor would be conducting this, rather than separate vendors. So I'm sorry if that wasn't clear. But – oh, okay. All right.

There are no – oh, sorry, Jamie, go ahead. You had your hand up.

JAMIE HEDLUND: One of the articles that I had assigned actually had a statistic for second-level domains that have been signed, and it was about 3%.

BRIAN AITCHISON: Right. And those numbers are out there, as Jamie indicated. So I'm not sure if we want to just limit ourselves to the registry level, that top level. If we want to be very technical about our scope, I think that we probably might have to. But just in terms of meaningfulness, I think that's an important considering. While it's fully deployed at the top level, at the second level it's 3%. So it might be something we can just mention.

Okay. It sounds like we think this is a pretty important one and amenable to quantitative inquiry, could use some help from a vendor.

Okay, great, prohibition of wildcarding. And I'll say this about the next two. Prohibition of wildcarding and removal of orphan glue records, they would both require qualitative approaches. There's been no complaints regarding each. There are complaint submission forms via the ICANN website. So there's no real quantitative data to play with. And as far as I know – my expertise is fairly limited – but I don't see how deployment of these two safeguards could be statistically correlated to a DNS abuse rate, but correct me if I'm wrong.

And the reason I'm grouping these two together is, in reading through the public comments and talking to experts and building the DNS Abuse Report, there doesn't seem to be much controversy around these. They're deployed. They're generally perceived as effective. In my early view, it doesn't seem like there's much to say about them, other than they're there and they seem to be working, and no one's really complaining about them. But I would love to get your feedback on that, if you see a better way to approach it, or more nuanced way.

CARLTON SAMUELS: I looked at the SSAC reports on this issue. And I was kind of challenged to find what has happened since the recommendation from the SSAC. So I would be at least interested to know where this is occurring or where tit occurred. There was one mentioned in the SSAC report, one registry operator that actually requested it and it was turned down. And it was a very cryptic reference. And it does pique my interest to see whether or not there were any other requests, or whether not the safeguard itself was instituted across registries and how broadly that

was put in. So for me personally, I want to know a little bit more about it.

BRIAN AITCHISON: And also, to be clear, I don't want to feel like we're skipping over or eliminating any. I think, at a minimum level, these will require some form of descriptive study, what's happening. But it might not necessarily require a vendor to conduct a six-month analysis of what's happening with wildcarding and orphan glue records. So you'll see in my bolded points here where I thought ICANN could take the lead role, as we have plenty of subject matter experts that we can talk to that can give us a good idea of what's happening, who know a lot about it, and address it that way, rather than through some sort of inferential way.

Carlton?

CARLTON SAMUELS: I'm glad you mentioned that, because in the SSAC reports, they identified some of the authors. And these guys, I'm sure, would be willing to assist. And they did put themselves in line for further assistance, as far as I read. I'll send it around.

BRIAN AITCHISON: So Thick WHOIS, this is obviously a contentious one, probably one we'll want to address pretty extensively, given the controversies surrounding it. The issues are how to measure the effectiveness of Thick WHOIS. And one second. There's a few possible ways to do it, and perhaps others. One is we have a WHOIS accuracy reporting system, which suffers from

some deficiencies in that it doesn't account for privacy and proxy services, among others. But it's a possible way to correlate. If there's a high accuracy rate among WHOIS, does that correlate in some way with some kind of TLD abuse rate? My basic point is that there's numbers there that we can play with and perhaps correlate.

Alternatively, or in addition, we could file this under a "perception of effectiveness" survey, where we survey. How well has Thick WHOIS worked for you? What are issues? These kinds of things. I view it as a safeguard that would require a vendor, or a vendor would be useful, to conduct a broad analysis. But that's my two cents on it.

And, Carlos, sorry, I'll turn it over to you.

CARLOS RAUL GUTIERREZ: Yes, this issue has not only another PDP, but the separate [AOAC] review. I agree with you, it's a hot potato. But this should not be our main worry, and we should avoid any overlap with the PDP and the other reviews. And so it is not that it's not hot and priority, but it's not our responsibility. Thank you.

BRIAN AITCHISON: Response to that? Laureen?

LAUREEN KAPIN: Yeah, I agree with you that it's a hot-button issue. It just strikes me, if you were trying to measure this, a lot of the entities that rely on the Thick WHOIS, of course, you're not going to get that from users, do you

care about Thick WHOIS? You're going to probably want to ask law enforcement and people doing investigative work, whether in the private sector or the government, about that issue. So I just wanted to make sure that gets out there. You really need to think about who your intended audience is, so to speak.

BRIAN AITCHISON: To respond to that, that's one reason why I think a vendor could be useful, because they would know how to segment, I would imagine, or have a capability to approach certain segments as we define them. So that's a useful point, and I've captured it.

Drew?

DREW BAGLEY: I'm glad you put "accuracy" in quotes, because that category, that's part and parcel as to why it's a controversial issue, is because something can be operationally accurate and have the proper syntax and be completely fake information, as far as everything else is concerned. And so, yeah, I agree with what was already said. It's already being looked at by other groups. So our value added on this would just be whatever correlations we could draw between this and the abuse data we're getting. And then from there, we might be able to say that even if the information is completely accurate under current definitions, that still means this much abuse is prevalent, or these safeguards requiring both operational and syntax validity do actually have a correlation with a lower rate of abuse, even if the definition of "accuracy" is controversial in and of itself.

BRIAN AITCHISON: Just add a caveat in there.

DREW BAGLEY: Yeah.

BRIAN AITCHISON: Okay.

DREW BAGLEY: I don't think we can draw a whole lot of useful information out of the data now.

BRIAN AITCHISON: Right. Jamie?

JAMIE HEDLUND: So two things, as I'm sure everyone is already aware. One is there is a lot of Thick WHOIS policy implementation work going on with respect to legacy gTLDs. And so there may be more useful comparative information in the future, when that's actually done. But the other thing is this will also be part of the WHOIS review, will likely be part of the WHOIS review team later on. If there is a candidate for something to fall off, this would seem to be one of them.

BRIAN AITCHISON: Okay, that's great. Okay. Anything else on Thick WHOIS? Okay.

So we have the centralization of zone file access. I would see this as qualitative kind of approach. I would see this as falling under a kind of survey of users of the CZDS, how effective they see it as being. I don't see any way to necessarily measure it quantitatively. Do we have any thoughts on that?

Jamie?

JAMIE HEDLUND: Sure, thanks. So this was also one of my reading assignments. There is a monthly report that ICANN keeps of credentials for the different zone files. And I'm not sure that those reports tell you anything in and of themselves, other than some generics are a lot more interesting, or have a lot more credentials, than brands, for example. So other than that, I don't know what… Other than there's an ease of use compared to what there was before, for legacy TLDs, my understanding is you have to negotiate a contract with each one each time you want to see something. You sign up once, you get credentials for the TLDs whose zone files you want to see, so long as you qualify. But other than that, I don't know. It's not clear to me what the data shows, other than some are more viewed than others.

BRIAN AITCHISON: Are these the access passwords, the zone file access passwords? Okay, thanks.

Drew?

DREW BAGLEY: Yeah, I was just going to say this falls under a category that would have interest, as far as safeguards, for cybersecurity researchers who would be looking at trying -- an external group who would be trying to do their own analyses on DNS abuse. And so I think this is only part of the puzzle. Historical WHOIS data, if that was public, that would fall into here too and whatnot. But this is a safeguard that, like journalism or anything else where other groups are adding value, and that's what can help it become a safeguard.

BRIAN AITCHISON: Right. And I think something similar, a similar segment could or should be targeted, as Laureen mentioned, with Thick WHOIS, if we were to conduct that. We want to talk to law enforcement, perhaps trademark holders, I know, use it a lot. So it could be useful for survey in a segmented, targeted survey. So, great.

UNIDENTIFIED FEMALE: So, Brian, I'm mindful of the time, because we're at 12:45. And I'm also mindful of wanting to hear the rest of your discussion. So I'm going to make a proposal that we have two choices. One is to continue with this discussion and have a shorter lunchbreak, or take a very quick break to run out, get grub, and come back and eat while we continue to listen to Brian. And maybe someone can bring Brian back some food so he doesn't faint away.

What's folks' preferences?

UNIDENTIFIED FEMALE: Can I just interrupt?

UNIDENTIFIED FEMALE: Yes.

UNIDENTIFIED FEMALE: Your lunch is way across the hotel, in the main dining room.

UNIDENTIFIED FEMALE: Oh, I didn't realize that. Thank you for that fact. I thought we were running out to get our own lunch. Yes, that's what we government workers do. What we lack in civilized, we make up for in efficiency.

So it sounds like then maybe we can take, if folks don't object, a little more time to hear what we can of Brian's presentation, and then have a slightly shorter lunch break. Apologies. Does that sound reasonable, folks? Yes? Yeah? Yeah, okay.

BRIAN AITCHISON: Sure. So we have until – how about we finish just these DNS abuse discussions and then break for lunch?

UNIDENTIFIED FEMALE: Sounds good.

BRIAN AITCHISON: Great. Okay. This is actually very useful. And so what I'm thinking of doing, once I consolidate all this, find some way to get all your feedback on actual decision points and where we feel like we want to head with these.

So anyway, moving right along, another safeguard is documented –

UNIDENTIFIED FEMALE: Brian?

BRIAN AITCHISON: Am I…

UNIDENTIFIED FEMALE: Sorry, just before you go on, I've been asked by a couple, is this document on the wiki somewhere? People are having a hard time reading the fine print from the screen. So is there a way you can send it around? Or is there a place where it already is?

BRIAN AITCHISON: I can send it around again. Should I just send it to the Safeguards Team?

UNIDENTIFIED FEMALE: Yeah.

BRIAN AITCHISON: Okay. One second. Let me just do that. Here we go, okay. So hopefully you all should be getting that.

Okay. Documented registry and registrar-level abuse contacts, I feel like this, again, lends itself to a qualitative approach. I called this a low bang for buck, in terms of hiring a vendor. It's a bit of a zero sum. Is it there or not? How deeply we want to delve and how effective having those contacts available is something that could lend itself to a survey; again, a segmented survey of, how useful have these been? Sort of ironic initial feedback has been that these are used by spam harvesters to send spam. So I'm not sure how we would want to approach this. So are there any thoughts on it? It's one of those hard-to-measure ones.

DREW BAGLEY: Is that data already collected by ICANN Compliance? Because that is a compliance issue.

BRIAN AITCHISON: It could be.

DREW BAGLEY: Because I think I've seen that in their reports. Yeah, can anyone else chime in on that? It might already exist.

BRIAN AITCHISON: Yeah.

UNIDENTIFIED FEMALE: Yeah. In looking at a lot of the Compliance data, they do monthly reports, yearly reports. And if there are complaints about a lack of an abuse point of contact or some other lack of publication of the contact information, that is information that ICANN Compliance collects, because it's one of their complaint categories.

DREW BAGLEY: Because the value in that would be if we saw that we had some measurement about the perception of who someone should contact if they saw DNS abuse and they perceived it wasn't the registry or the registrar, because there was no abuse contact listed. I guess there's stuff we could slice that way. But I agree with what you're saying about the bang for the buck, and I think the data already exists, so we could incorporate that into a report for contextual purposes too.

BRIAN AITCHISON: Right. I tend to think it would be fairly descriptive. There might be some outstanding issues, but, yeah, okay.

Okay. Any other comments on that safeguard?

Okay, not seeing any, here's another difficult one, the expedited registry security request process. There haven't been many instances of its use. I believe it's something on the order of two or three. Not only that; it tends to be quite sensitive as to what it was responding to and how it was responded to. Just for security reasons, there's not too much we can delve into. So I've flagged that as something perhaps ICANN can speak with subject matter experts about. We have people in our

security community who know more about it. Again, a vendor could conduct surveys as well. A vendor could fit into any of these, if we wanted to force fit it.

But do we have any thoughts? Do we see a higher bang for buck anywhere on this one?

I'm going to take that as a no. Great. And now, the final one, which as a safeguard doesn't really exist, but if you look at the public comments it generated perhaps the most substantive and strong responses. I'm not sure how we can measure a safeguard that doesn't formally exist. My initial idea was to review the public comments, interview some SMEs, and provide just a description of why the safeguard wasn't adopted.

There are several examples of independent registry operators implementing their own kind of HSE-type program. So I think that falls into an ICANN descriptive kind of study. But I'm happy to know what you all think about that one, if there's any way to measure it or discuss it.

Yeah, yeah, there's not really much to do with it, so.

UNIDENTIFIED MALE: [inaudible]

BRIAN AITCHISON: Okay, so I'll flag that as low. Okay, well, that's our nine DNS abuse safeguards. I hope this is useful. Are you finding this is giving us some direction? Is there anything I could improve or answer better?

UNIDENTIFIED FEMALE: First of all, I think this is enormously useful, and I appreciate you putting together this document. I think the robustness of our discussion is somewhat hobbled by the fact that people haven't had a chance to digest this and look at it and, in some cases, can't see it very well. So what I would suggest is that we give people an opportunity to digest it a little more, and then revisit this so that we can revisit it in the short term, so that we can complete our discussion on this. Because I think it's a very useful exercise, but I just think people haven't had a chance to digest it enough.

Okay, so extra-special thanks to Brian for all of this hard work and this effort, which I think will be very useful. So we have a lunch break. Hopefully we'll have a guide to take us to said destination for our vittles. And then we come back here at 1:45 to continue our discussion.

Questions? Carlos, yes.

CARLTON SAMUELS: [inaudible]

UNIDENTIFIED FEMALE: Yeah, that's actually – I had the same thought. Can anyone… Somebody will be in here?

UNIDENTIFIED MALE: [inaudible]

UNIDENTIFIED FEMALE: So [caution] yourselves accordingly. Any other questions?

[END OF TRANSCRIPTION]