
COPENHAGEN – Engagement Session with the Competition, Consumer Trust, and Consumer Choice Review Team
Sunday, March 12, 2017 – 09:00 to 10:30 CET
ICANN58 | Copenhagen, Denmark

JONATHAN ZUCK: ...conclusions. There's a famous Danish saying that there's no cow on the ice, which means that there's nothing dramatic to report. The improvement in competition, consumer choice, and the adoption of safeguards, is there. There's indicators that there's been an increase in competition. There's obviously an increase in choice. And adoption of safeguards has been largely successful.

I think if there's any elephant in the room, it's about data itself and the availability of data within the ICANN ecosystem. And that's something we find ourselves talking about a lot and it happens across reports and across reviews, that data was a hindrance to doing a more thorough evaluation of competition and safeguards. And so you'll see us talking about that a lot. We generated a lot of primary data but I think there's more data to be got, and at the heart of our recommendations is expanding the availability of data inside of ICANN.

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.

We plan on doing a webinar to go through this in some detail for people that are working on submitting comments. We weren't able to nail down a date but it's likely to be one of the last three days of March – the 28th, 29th, and 30th – one of those days we'll do a couple of webinars to help people get up to speed and answer questions. So as you start reading this report this afternoon, it gives you a few days to generate questions.

We really do welcome your input and feedback, both today, over the course of the week, and hopefully also as part of the public comment period. The public comment period is now open and goes through the 27th of April, so there's a pretty significant amount of time to get through the report, to make observations about it, and help us to make it the best report that it can be. The easiest place is [comments.cct-rt-draft-report](#). Never mind. We'll put that on the [cct.wiki](#) page to make it easy to find.

Alright. Next slide please.

Key findings – on balance, expansion of the DNS marketplace has demonstrated increased competition and consumer choice, and it's been somewhat successful in mitigating its impact on consumer trust, and rights and trademark protection – it appears as though the safeguards have been in place for rights protection have been largely successful as well.

The caveats – the New gTLD Program is really just a kind of a good start. It's really early in the program. It's difficult to make assessments. Historically it's been around year three that a new TLD has really hit its stride, and so it's difficult in the short period of time that we looked at – which is the first year of the program during which TLDs were being allocated – that we were studying. So it is just a short period of time.

A number of policy issues should be addressed prior to any further expansion of the TLD space. And as you see, when we talk about our recommendations, there are some that we've listed as prerequisites to moving forward with the New gTLD Program.

So incomplete data limits more comprehensive analysis – you'll hear me mention data a lot.

Next slide.

If you look at our recommendations, the new version of the reviews as modified by the CCWG and put into the Bylaws – so it used to be part of the Affirmation of Commitments and now it's a requirement of the Bylaws – there's a new requirement that any recommendations by a review team need to specify whether or not they need to be completed prior to a continuation of the program, and so we came up with a few more categories than

that to try to get at the idea of timelines for implementation because there are 50 of them we hope that they're less intimidating from an implementation standpoint because they're kind of spread out over a timeline.

The ones that are labeled as “prerequisite” need to get done prior to any further expansion of the new gTLD space. If they're “high priority,” then we want to see them get done within the next 18 months, “medium priority” within 26 months of the final report, and then “low priority” is really about having enough data to do this review the next time and so those things in that category – there's eight of them – are data collection, etc. that needs to happen prior to the next CCT Review.

Next slide.

Just looking at the recommendations in sort of a snapshot form, you can see that there's a breakdown between data analysis recommendations, competition, choice, trust, safeguards, and the application and evaluation process. So there's a pretty good distribution of the recommendations, butt “safeguards” had the broadest number of things that's being covered which you'll see today and so that has the most recommendations associated with it.

Next slide.

Briefly on data collection – we tried very hard to make this a data-driven effort and to use data to assess the state of the New gTLD Program and the DNS marketplace post the introduction of new gTLDs. We did a consumer survey of end users. We did a survey of registrants to get a sense of their trust and the impact of the trust and preferences of consumers in the DNS. An economic study was done by Analysis Group to look at the competitive effects of the new gTLDs. There was an applicant survey where we asked the folks that applied for new strings what the process was like to participate in the application process.

And then in an effort to better understand the lack of applications from the Global South, we actually conducted a kind of a cohort survey – which is a little bit meta – but the idea was to identify similar organizations to the Global South to the ones that did apply from the Global North and ask them why they didn't apply. So it's an interesting study to read as well.

There was also a parking rates study that was just completed, and there are two studies that are still out. One is on safeguards against DNS abuse and the other is a survey of INTA members about the impact on rights holders.

So there's still additional data that we feel like we were missing and would like to have had – pricing and wholesale retail and

secondary global and regional data, etc. In order to really understand competition you want a better understanding of prices.

“Competition analysis in the standpoint of substitution behavior” – A better understanding of user behaviors. When they register a new domain, is it replacing an old one or is it just a supplement?

“Tracking of programs intended to facilitate applications” – This is part of the Applicant Support Program that we’ll talk about a little bit.

“Subject matter of complaints” – Looking for much more granular data collection by the Compliance Department in order to better understand the impacts of [individual] safeguards.

And then finally, “DNS abuse rates in legacy versus new gTLDs.”

So there’s a lot of data still to be collected and understood to really understand the impact of the New gTLD Program.

Next slide.

I would say that one of our most fundamental recommendations is in a way, more general and generic because it does come up so often, which is: “To formalize and promote ongoing data collection inside of ICANN.” And so we’re even recommending

that there be a dedicated data scientist as part of ICANN so that working groups, strategic plans, etc. all have the benefit of identifying up front what data would make sense to collect so that success metrics can be defined around that data, etc. And that's something the organization has needed desperately and we just felt it very acutely in our review. So this particular recommendation is meant not just for CCT purposes, but for a better ICANN generally.

Next slide.

Any questions? That's just the overview. We'll go over things in more detail, but if you don't have any questions now then I'm going to hand the microphone over to Jordyn to talk about competition and consumer choice.

JORDYN BUCHANAN:

I'm Jordyn Buchanan. I Chaired the Subteam on Competition and Consumer Choice – and if we move on to the next slide we can talk about what we've been up to.

Although it seems like we've been working at this for years, it's only been about 14 months since we kicked off. In fact though, the base groundwork in order to do our competition and consumer choice analysis has been going on for years. We relied quite a bit on previous work done by Analysis Group in the form

of their economic study as well as Nielsen in their consumer and registrant surveys in order to inform our work and collect data relevant to the issues of competition and consumer choice.

One of the things that we had to do once we started to look at the raw data and even the aggregated data that Analysis Group had put together was, on the subject of competition you really need to think about competition within a specific market. It's been a challenge, I think, for us to think about what exactly the market definition that we should use would be. And so we ended up, to a certain extent, not making a hard choice and proposed various different definitions of relevant markets.

For example, you might think that the market in which a new registry operator, a new gTLD, would compete, would be against the legacy gTLDs or all of the gTLDs, or someone else might propose that the marketplace in which the competition is occurring is actually not just gTLDs but they also compete with ccTLDs, or someone else might propose that it's actually within this particular TLD is in say Spanish or something like that so it's really the cohort of Spanish language TLDs and ccTLDs that are relevant. And so as opposed to trying to get to the one true definition of "market," we proposed various different definitions and performed calculations on market shares and concentration for each of these different market definitions.

The key to understanding the initial trend at least, was to look at the measures basically at the very end of 2013 which is right before the new gTLDs started to get delegated, and try to do comparisons over time but particularly focusing on March, 2016, as a comparative date which is one we had good data for this initial report.

We do believe that we'll have data as of December, 2016, in the final report that we expect to come out later in the year. So one of the things that we'll be doing is to revise many of our calculations with more recent data.

Next slide please.

As Jonathan said, we generally observe positive signals relevant to competition at this phase. For example, one of the statistics that I think is easiest to wrap our heads around is that if you look as of March '16 and starting at the end of 2013 date, about half of the growth in gTLD registrations or half of the growth in the registered gTLD space, was new gTLDs versus the other half being legacy gTLDs over that period of end of 2013 through March of 2016. And so the new gTLDs as a whole represented about as much growth in the gTLD space as the legacy gTLDs.

Sort of, I think coincidental, but interesting observation that is if you throw ccTLDs into the mix, instead of being 50/50 as was the

case with legacy and new gTLDs, it becomes about a third, a third, a third – a third of the growth in registrations went to ccTLDs, a third to legacy gTLDs, and a third to new gTLDs. So if you think about those three being chunks of TLDs in aggregate, the new gTLDs sort of roughly held their own against the other two chunks.

The other thing that was important when we think about competition is to think about how concentrated the marketplace is. And this is something that competition authorities look at a lot. They have some fairly technical measurements that they use in order to understand concentration. And the concentration is roughly what share of the market is aggregated in the biggest players? Historically speaking, the gTLD market in particular has been quite concentrated. As probably everyone here is aware, the .com registry is by far the largest of the gTLD registries and the next largest gTLD - .net – is actually run by the same registry operator and so when we think about concentration we look at which one registry operator, how much market share do they have? Historically that's been very concentrated in the biggest registry operator.

However, what we've observed with the introduction of new gTLDs is that the standard measures of market concentration

have improved substantially, and by this I mean they generally showed decreases in the numbers that we use. And so there's various ways of looking at it, but the HHI measurement which is the one that the U.S. Department of Justice tends to go to, declined by over 1,000 points between 2013 and 2016. That number is still quite high and above the point at which I think anti-trust regulators at least consider the HHI number to be interesting, but it has gone down considerably which we view as a positive trend.

I'll note just as an aside – it's not on the slide – but competition between registrars or concentration in the registrar marketplace was already quite low, and that remained the case throughout the period that we studied. And the other factor that we looked at was competition between backend providers for registry services, and that also showed improvement along with the competition between registry operators.

Next slide please.

Usually what you would expect to see if market concentration was decreasing is that prices would improve, prices would decrease. We did not observe that. And it's mostly because we don't think we were able to observe that largely because we were unable to determine what changes of price were happening in the legacy gTLD places. And there's really two

factors involved here. Number one is, unlike new gTLDs which are required to cooperate in economic studies, legacy gTLDs don't have that obligation in their contract with ICANN. So in many cases we were unable to get the data that we had hoped for from the legacy gTLD operators to see if there was a trend. But it probably wouldn't have mattered anyways because the legacy gTLDs all are subject to price caps, and in reality the new gTLDs tend – one thing we did [find] is that on average the new gTLDs tend to charge more than the legacy gTLD price caps. So the impression that you might get is that the fair market price for a domain registration at a wholesale level might be somewhat higher than the price caps that are currently in place for legacy gTLDs, and therefore even if the new gTLDs were putting a downward pressure on price, the price that the legacy gTLD operators would want to charge would still be higher than their caps and so the caps are really constraining the marketplace analysis at this point.

The last big takeaway from the competition side is that even though on aggregate the new gTLDs I think have competed quite effectively with the legacy gTLDs, most of them are quite small. Almost three quarters of the new gTLDs that we took a look at have less than 10,000 registrations, and almost all of them – 90% – have less than 50,000 registrations. And that's even after

backing out all of the .brands and single registrant TLDs. So this is of the TLDs that are held out to the marketplace for sale. So they're quite small.

We did take a look at this, and it's possible that even though they're quite small, that the minimum viable scale for a registry is also quite small because registries can rely on registrars, for example, to do their distribution and backend service providers to do their technical operation. So it's possible that they don't need that much in the way of their own operations in order to stay in business. This is something that we'll have to keep an eye on in future reviews to see whether the scale can be maintained over time.

Next slide please.

So moving on from the topic of competition, we also took a look at consumer choice. And there was a sort of obvious result which is that the new gTLDs give consumers more choices. And this is true for both registrants who now have more types of domains to choose from – new languages, new character sets, geographic identifiers, and specialized categories – and it's also true for consumer end users – people navigating to those domains – who have more options and can get greater specificity in identifying the domain name that they're navigating to. That was I think a

fairly obvious conclusion. But if we move on to the next slide we'll see some of the other elements that we took a look at.

One of the things that we tried to understand was whether along with the options for consumers and for registrants in particular, there's cost to them potentially. For example, if instead of registering a new name because they felt like it enhanced their brand or because the domain itself was attractive and helped them reach new customers, in some cases they might feel obliged to register in these names – we call this “defensive registration.” And if they were registering because they felt obliged to, that's just the cost without a benefit associated with them.

Our research through the Nielsen survey and other data points that we have, showed that while some registrants did feel that they were obliged to and were registering with defensive objectives, we did still see a large number of registrants that were registering in order to broaden the appeal of their offerings or make other productive use of the gTLDs.

We took a very specific look at the defensive costs for trademark holders in particular, and this largely keys off work that Analysis Group did as part of the Trademark Clearinghouse Review. What we found through that analysis was that in general, ICANN's put together a large number of Rights Protection Mechanisms that

are designed to minimize the cost of the introduction of new gTLDs on trademark holders. And in general, those seem to be working pretty well. Most of the trademark holders that had registered in the legacy gTLDs either had registered only a small number of domains in the new gTLDs or they hadn't registered at all in the new gTLDs. However, we did see a small number – I think 4% – of trademark holders in the Analysis Group study had registered over 100 versions of their trademark in the new gTLDs, and we found one in particular had registered in over 400 of the new gTLDs. So it's possible that there's something of a bimodal distribution here and a small number of trademark holders are really bearing the brunt of the cost of the defensive registrations in the New gTLD Program.

So we make a specific recommendation here and ask that the Policy Development Process consider whether there's a mechanism to minimize the cost of defensive registration for those brands that are currently registering many, many, instances of their domain across many of the gTLDs.

Next slide please.

We also took a quick look at registry policies. One aspect of consumer choice [that] might exist is that you might choose a TLD based on what policies that TLD has. We did find that there was a variety of different policies that had rolled out there and

we took a very close look at privacy policies in particular. One thing that we did find is that not very many of the top 30 new gTLDs had very clear statements in their policies about sharing or selling registrants' personal data and so we make a recommendation here as well to have clearer statements and a little bit more regulation around the collection and sharing of data by registries.

Next slide. I think I might be at the end.

So beyond what we've done so far we have a few more months between this and the issuance of our final report. We plan to do a few things, as I mentioned earlier. The first thing that we plan to do is update our calculations with more recent data. The second thing that we plan to do is take a look at the topic of domain parking. I think there's been some speculation that speculation is driving much of the registrations in the new gTLDs and we do indeed see higher rates of what's called "parking" or registering a domain and then not using it or putting up ads on it in the new gTLDs than in the legacy gTLDs, and there's a theory that this may result in lower average renewal rates. If that were true, then our estimates of the long-term market penetration for the new gTLDs might be overstated. So we're going to take a closer look at the parking data and also revise some of our

calculations considering what they might look like if you backed out some of the domains that are currently parked.

And then the last thing we're hoping to do is to get a handle on more specific market segments is to look at data collected as part of the LAC study which actually has WHOIS information associated with individual registrations, and try to do some country-specific market shares as opposed to just looking at global market shares like we've done so far, which may help give us some insights about whether we see differences in behavior in different specific geographic markets.

That's where we are on competition and consumer choice so far. Like Jonathan said previously, generally positive signals, but I'm happy to answer any questions at this point that people may have about the data or our findings. Thanks.

JONATHAN ZUCK: Be sure and introduce yourself [too].

OWEN DELONG: Owen DeLong, Akamai. You mentioned that you're bemoaning the lack of data from the legacy gTLDs on pricing and you also mentioned that they're currently subject to price caps which may actually give them somewhat of an economic incentive to

be used versus the new gTLDs. I'm wondering if there's been any consideration to a recommendation that the price caps be removed in exchange for a contractual provision that they also cooperate with economic studies.

JORDYN BUCHANAN:

There is a recommendation that the legacy gTLDs participate in economic studies in the future. So we got halfway there. In regards to removing the price cap for legacy gTLDs, we don't quite reach that point. I think we do make an observation that the competition with new gTLDs may be robust enough that the legacy gTLD price caps may no longer be required, but we don't get to the point of recommending that the price caps be removed largely because that's not really in our charter. We're chartered to look at the effects of the new gTLD program as opposed to looking at the legacy gTLD competition dynamics in particular. And there's a thorny set of issues around I think .com in particular where the U.S. government's actually involved in setting that price cap.

So I don't think it would be quite within our mandate to get there, but we do make observations at least that the competition that's been provided by new gTLDs may make it worth taking another look at that topic at least.

Any other questions?

UNIDENTIFIED FEMALE: We have a remote participant question: “Can you please comment on the recent articles regarding significant price increases on some new gTLDs?”

JORDYN BUCHANAN: Certainly. We actually in the report, when we look at this issue of small gTLDs, there’s a number – we identified five – possible outcomes as a result of the small scale of some of the gTLDs. These range from, “actually it may be fine because the scale’s sufficiently small,” to, “maybe someone will buy them,” “maybe they’ll go out of business,” But one of the things that we identified might happen is that prices might increase in order for a niche provider to identify a particular market in order to be able to serve that market and still stay in business. So I think I guess I would say that it is not unexpected that price increases happen in some of the smaller gTLDs. I don’t think we made any speculation or have a reaction to the magnitude of those increases, but the phenomenon of seeing increased prices in some of the smaller gTLDs is one of the options that we thought might be expected as a result of the small scale of some of the new gTLDs.

UNIDENTIFIED FEMALE: Thank you, Jordyn. And that question did come from Vicky Sheckler. Thank you.

JORDYN BUCHANAN: Great.

JONATHAN ZUCK: Any other questions?

JORDYN BUCHANAN: Alright. I will turn it over to Laureen, I think.

LAUREEN KAPIN: Good morning. Our subteam focused on safeguards and consumer trust issues, and you're going to be hearing from me as well as my colleagues Drew Bagley and David Taylor.

Next slide please.

We had some very key findings regarding consumer trust and safeguards. I'm first going to talk with you about consumer trust and then I'm going to pass the baton. But as an overview for consumer trust, we didn't see any reduction of consumer trust in the DNS overall after the introduction of the new gTLDs. We did

see some links between the levels of consumer trust and various factors. Certain factors were particularly important to the public when they considered what gTLDs they are going to trust, and those factors were familiarity – which includes awareness of a gTLD and reputation of the gTLD. What we saw was, the more familiar the public was with the gTLD the more they tended to trust it. And also the other concerns that were evidenced from the Nielsen studies which focused on consumer end users and registrants were security. Did the public have a perception that a particular gTLD instituted security measures, for example, to prevent them from getting ripped off, to prevent their personal information from being shared without their knowledge, things like that?

So these two issues – familiarity and security – really had an impact on consumer trust vis-à-vis a particular gTLD. But we found that even more work could be done to study this issue because the studies that were done really focused in large part subjectively – how did the survey participants feel about a certain gTLD? Did they trust it? – but we think there could be more work done to study objective factors. Why do consumers choose to visit and become familiar with certain gTLDs more than others? What sort of behavior indicates trust of a gTLD, i.e. are you willing to share information with that gTLD? Are you

willing to engage in transactions with that gTLD, for example. So we think there's more work to be done there.

At a high level, regarding safeguards there were a large number of safeguards implemented via contracts for the New gTLD Program, safeguards that didn't exist before. So that in and of itself is an improvement.

But like the consumer trust topic, we still think there's more significant information that could be gathered regarding these safeguards which were intended to protect the public but we also would like to know whether that's actually happening.

And then we also saw a real need to collect information specifically focused on DNS abuse because many of the safeguards actually were targeted at preventing DNS abuse, and we'd like to see whether that's actually happening. You're going to be hearing from my colleague Drew Bagley about a study that this group helped to launch which is to really start collecting information on DNS abuse, at least looking at correlations and comparing legacy gTLDs and new gTLDs and trying to correlate what factors might link to DNS abuse.

Finally, one of our big recommendations that I'll be telling you more about is trying to get even more transparency about the complaints filed at ICANN in terms of what is the specific subject

matter of the complaints so we can better track what's happening in the New gTLD Program and the Legacy gTLD Program by looking at complaints and being able to find out what subject areas these complaints are really talking about.

Next slide.

In terms of consumer trust, what methodology did we use? I'm going to focus now on consumer trust. You'll be hearing a little bit more from me about safeguards later on. But for now, I have my consumer trust lens on.

The studies that we really focused on to grapple with this issue were commissioned before this review team even formed. And there were two studies that repeated in two waves – in 2015 and 2016. They were commissioned by ICANN, and Nielsen performed these studies. One focused on consumer end users and the other focused on registrants. And they were rather lengthy surveys that asked a lot of questions about a lot of topics, but very particularly focused on consumer trust issues in certain parts.

So what were our key findings there? Those I've already touched on – that familiarity and security really drive the public's trust of new gTLDs and gTLDs in general. But the new gTLDs haven't been around as long. People don't know about them as much.

So not surprisingly, those studies indicated that the public only trusts new gTLDs about half as much as they trust legacy gTLDs.

One of our recommendations is that these studies get repeated. These are a baseline. So we may find as the public becomes more familiar with new gTLDs that trust levels may change.

One of the other primary findings we found is that the public tended to find a relationship between registration restrictions on who can buy a gTLD were there restrictions, for example, were there certain screenings to make sure that entities that bought a gTLD in a regulated industry perhaps like a pharmacy or a bank, if there were restrictions there was a correlation between those restrictions and trust. The public tended to trust gTLDs that had restrictions.

But overall in terms of big picture, we did not find that trust in the DNS overall diminished after the expansion of the New gTLD Program.

Next slide.

So what were some of our key recommendations? And I want everyone to realize this is just sort of a highlights trailer reel, but if you want to see the whole movie which has great content then I want to urge you to read the whole report. But for some highlights about our recommendations we really would like to

collect some more data. So as I talked about, we'd like to really see which new gTLDs were visited and why and how users' actual behavior relates to trust. Those we feel would give us a little more objective information.

We also found as a result of the surveys that the public had certain expectations. One expectation was that consumers expect a relationship between the name of the gTLD and its content. So for example, if there's a .photos gTLD, the public expects that when they go to domains within that gTLD they are going to see things that relate to photos, not typewriters. We also found that there's an expectation about implied messages of trust conveyed by a name. So if I'm seeing a .accountant gTLD, my trust that when I'm dealing with folks behind that gTLD that they're going to treat my financial account information in a secure way. And also the public has expectations that when they provide sensitive information, for example, health information or financial information, that that information's going to be treated securely.

So one of our recommendations is to create incentives to encourage gTLD registries to meet user expectations about these issues, to encourage gTLD registries to raise the bar and think about what the public expects.

And finally, as I already mentioned, we have these studies that have provided a lot of useful information but they're a baseline and we want to see how things change as the New gTLD Program evolves and matures, so we'd like to see these studies repeated.

I'm going to perhaps take a little mini pause before I pass the baton over to my colleague Drew to talk a little bit about DNS abuse. So if folks have questions about consumer trust issues, now would be a good time to ask me.

OWEN DELONG:

On Recommendation #13, I'm curious as to a little bit of the technicality of how you go about measuring that given that most people are primarily visiting some form of SLD within the gTLD and that once they've cached the gTLD server information, they may not come back to the root name servers to "visit" the gTLD again to go to the next SLD within the same gTLD. So you may not necessarily have a particularly accurate way of getting at how people are visiting gTLDs due to caching.

LAUREEN KAPIN:

You raise an interesting point and a lot of technical issues which I'll confess is not my particular area of expertise. But this is something that the review team and the folks who stay on to

help implement would actually work with whoever is commissioned to do these studies. So we would have to come up with a way to measure what we're interested in, and I think that is an interesting and very specific comment you're raising about the technicalities and I would encourage you to please submit a comment to the review team during the public comment period so we can take that into account because we are trying to be accurate and work with the folks who perform these studies to actually measure what we're intending to measure. So I appreciate that.

I think I saw someone else who had a question.

JONATHAN ZUCK: I may just add to that I guess a little bit.

LAUREEN KAPIN: Sure.

JONATHAN ZUCK: Just to clarify, this is about in some measure Recommendation #13 is about improving on the surveys that we did and asking different types of questions of people. And so some of that is less about the technicalities of tracking their behavior and more about asking them more specific questions. The survey was

somewhat handicapped by the fact that most of the people surveyed didn't even know there were new gTLDs. And so that meant that there hadn't been a dramatic impact on trust one way or another associated with the New gTLD Program. It meant that when we asked questions of trust etc., what we were getting back was not necessarily the reality of their experience with these new gTLDs but what their inclinations were. In other words, when I recognize the word, it feels familiar to me and therefore I'm more likely to trust it. And if I feel like the word is a semantic hint at what the content will be on that TLD, I'm more likely to trust it.

So in a way, what we got were some interesting harbingers for the future of consumer trust and a greatly expanded name space, more so than a reflection of the reality of today where most people simply aren't truly aware or have visited these sites. So those things will change over time as well. And so asking questions about their individual behavior in terms of where do they share their credit card information, etc. might be more objective criteria for trust than simply asking them, "Do you trust such and such a TLD?"

I think that's a little bit of where we were going with that. I guess Jordyn wants to add to [inaudible].

JORDYN BUCHANAN: Yeah, I would just add briefly that I think the intervention's a good one and that looking at DNS logs in particular might not be the right way to approach this study, so to Laureen's point, your feedback is appreciated but there are – I know Nielsen in particular and other firms – do have client instrumentation that might be a better choice where we can actually see what sites people are going to when they've opted into sharing that data. Thanks.

[DON]: [DON] [inaudible], .rs. My question is regarding the – you were speaking about trust in the DNS name space altogether, but I understand the studies were related only to gTLDs because this was the subject of your work. But in future, do you plan to do it in a wider way also to see about the trust in ccTLDs and maybe to compare the level of consumers' trust in all gTLDs, legacy, new ones, and ccTLDs? I think that will give us a bit broader picture. Thanks.

LAUREEN KAPIN: I can't speak for what future review teams are going to do. This is a general recommendation to 13 to try and get a more objective measure and 15 to repeat these studies. But you do raise a very significant point. Legacy gTLDs and new gTLDs are

only part of the picture and ccTLDs form another significant part. So again, you're going to hear this a lot but we're encouraging you to give us that feedback so that when we have this pause in time between our preliminary report and our final report, we can consider what the community is most interested in. So I appreciate your comment.

JONATHAN ZUCK:

And there were a few general questions, too, about trust in the DNS industry itself and that remains incredibly high despite what we all learn at these meetings. So I think that there's good consumer trust overall in the DNS in addition to just the gTLDs.

TIM SMITH:

Hi. Tim Smith from the Canadian International Pharmacy Association. You may have already answered this in your previous responses, but when you talk about a higher level of trust from gTLDs that require some screening for eligibility – and you mentioned bank and pharmacy – were those specific examples that you used in the survey?

LAUREEN KAPIN:

I'm trying to think whether bank and pharmacy were used in the survey, and I think that they may have been and I think also

.builder was used, all by way of examples of regulated gTLDs. The survey didn't say, "These are more trustworthy." That's not how the questions were phrased. The questions were phrased in terms of what the public's expectation would be in terms of restrictions for gTLDs that were in regulated industries.

One question that I remember specifically is, "Would you expect, for example, a .builder gTLD – someone who owns a domain in that gTLD – to have some sort of credential that they in fact are a builder or a general contractor?" That's the way the question was phrased. It was phrased in terms of consumer expectations. The questions were not phrased in terms of, "These are more trustworthy gTLDs than others." The questions were trying to get at what does the public expect when it is dealing with a gTLD in a regulated industry or where there's a perception that sensitive information is going to be shared because of the nature of the subject matter of that gTLD.

Unless we have questions from the chat, I'm going to pass the baton over to Drew. Are we good? Good.

DREW BAGLEY:

Thanks, Lauren. And thank you to all of you for being here on a Sunday morning.

Next slide please.

So one of our mandates dealt with DNS abuse because we were tasked with looking at malicious abuse issues related to the implementation of the New gTLD Program as well as the effectiveness of the safeguards that were implemented as part of the New gTLD Program. And these safeguards were developed with regard to several risks that were identified at the time that the New gTLD Program had been proposed. And so what we did was, we took a look at all of the safeguards including nine technical safeguards that were proposed and broader safeguards that related not always directly to DNS abuse but to the general concept of keeping the new gTLDs from having a disruptive impact for the risks identified several years ago.

And so as part of our methodology, what we did was we took a look at the implementation and compliance issues associated with each safeguard and saw whether or not, to the extent that we had data, it looked like the safeguards had been fully implemented by the new gTLD operators, whether there were complaints specific to their implementation of a safeguard, or a violation of some sort of safeguard. That's what we've done to create the framework of looking at these issues.

However, the biggest and most data-rich part of all of this will be this comprehensive DNS abuse study that we have commissioned that both Jonathan and Lauren have

mentioned, for which we have a vendor that is conducting the first analysis that will actually look at – rather than just doing a sampling – but look at all of the domain names registered in legacy TLDs as well as new gTLDs for the past several years, and compare different abuse rates based on the data feeds they're able to get dealing with phishing, botnet command and control, malware hosting, and some other issues.

So what we'll be able to see from that is we'll be able to do a comparative analysis to see if there is a difference between the rates of abuse in legacy gTLDs versus those in new gTLDs, and then hopefully see what correlations we might be able to determine along with safeguard implementation, as Lauren alluded to, to see if in fact the safeguards may have been an effective force in preventing and mitigating some of these abuse issues, if perhaps they could be improved upon, or if perhaps there was no change – no detectable, discernable, change – whatsoever.

Next slide please.

So from our initial findings, we found with regards to the nine technical safeguards at least – and those you can look at them in the report – but for example, deal with prohibiting orphaned glue records, ensuring that DNSSEC is implemented by the registry operators, and so on. We found actually widespread

implementation of those safeguards, and we also took a look at the number of compliance complaints for each of those categories but that's something from now with the draft report until we have the final report we plan on spending time with ICANN Compliance to see how we can get some more nuanced details about what may have been in those complaints. Nonetheless, thus far there's nothing that we've seen to lead us to believe there's any widespread problem with any of these as far as implementation is concerned.

And then as far as DNS abuse as a whole, looking at existing literature, as I mentioned there had been no comprehensive analysis of DNS abuse in legacy versus new gTLDs but we did find some existing research by other entities that at least did a sampling of DNS abuse.

And so from those studies from groups such as APWG, we've seen that perhaps there could be perhaps overall more abuse occurs in legacy gTLDs such as .com, but that on a percentage basis perhaps there are higher rates of abuse in some new gTLDs compared to legacy gTLDs when you're actually looking at the number of registrations versus the number of abusive registrations in a particular zone.

With that said, since these have only been snapshots, this really hasn't given us enough data to draw a definitive conclusion one

way or another, and so that's why we will be waiting for the results of this greater DNS abuse study that we've commissioned before drawing any sort of conclusions about that.

Next slide please.

And so thus far, we do not have many recommendations because we're waiting for the data, and so our only recommendation that we have so far is that this DNS abuse study that we've commissioned, we believe that this should be something that even if the methodology is improved upon or expanded going forward, we believe it's something that should occur regularly to help give a snapshot of what's going on in the new gTLDs versus the legacy gTLDs and believe that, as Jonathan has mentioned, that along with a lot of the other data we believe should exist in the ecosystem in which we all operate, we believe that this specifically would be very helpful for our successor review team as well as for other review teams such as the SSR2 and for just policy making within ICANN as a whole so that we're able to look at our study as the baseline and look at improvements or perhaps decreases in effectiveness over time with regard to the safeguards and DNS abuse.

That's where we're at specifically with DNS abuse so far. So for that, you're going to see a lot more to come in the coming months and by our final report, but I encourage everyone to take

a look at the draft report and give us feedback on this and I also encourage everyone to attend Tuesday’s session in which the vendors chosen for this study will actually speak specifically to their methodology in detail, and I think that’s a very unique opportunity to provide them with some feedback. And also what they’re looking for from the community, I know, is as many data feeds as possible. So to the extent that you know of good sources of abusive data feeds that they could potentially use in doing their statistical analysis that would be greatly appreciated so that as much data as exists can be relied upon and so that they could really have a statistically sound analysis.

Thank you. And I will entertain any questions related specifically to this topic. Otherwise, I will pass the baton.

Are there any questions?

UNIDENTIFIED FEMALE: We have a remote question, Drew, from John McCormack, HosterStats.com: “Does DNS abuse definition also cover compromised websites? Link injections seem to be the most widespread compromise at the moment.”

DREW BAGLEY: The vendor is absolutely making a distinction and looking at compromised domain names, not necessarily looking specifically at hosting. So if the domain name's been reported as being malicious and then what they're doing is they're looking at – and I will let them on Tuesday speak to the exact details of their methodology – but they are looking at how long a domain name was registered before it went bad, making distinctions based off of that as to a determination as to whether something was likely a legitimate registration that was later compromised versus something that was registered with a malicious intent to begin with. But I will certainly defer to them to go into more detail on Tuesday on that. But yes, in short, that issue is being looked at and considered.

UNIDENTIFIED FEMALE: Thank you.

DREW BAGLEY: So if there are no other questions, then I believe David Taylor – no, never mind. Laureen is up next. So if you're a David Taylor fan you have to wait even longer.

LAUREEN KAPIN: Always the best for last.

I want to talk a little bit about safeguards. You may have noticed that safeguards actually generated the largest number of recommendations. I'm not going to talk about all of them. But all of them are in our report. But I wanted to give you a flavor for our approach and highlight some of the recommendations we've made in certain areas.

In looking at the safeguards that were implemented as part of the New gTLD Program, most of these were implemented via the structure of contracts between ICANN, registries, registrars, and registrants. We looked at the goals of the particular safeguard at issue, how the safeguard was implemented and enforced, and then any issues that identified regarding particular safeguards.

I'm going to highlight some of our observations regarding safeguards related to WHOIS and sensitive regulated and highly regulated strings.

Next slide.

For the WHOIS related safeguards, the stated goal there is to enhance abuse prevention and mitigate abuses. And some of our findings in this area was that the largest category of complaints that ICANN Compliance received is in the WHOIS area – complaints regarding registrars. And in response to concerns about the accuracy of the WHOIS information being

collected, ICANN instituted a project, a WHOIS Accuracy Reporting System that was originally structured in three phases – syntax accuracy is the WHOIS information in the form is supposed to be, is it in the right form that an address should be or an e-mail should be – syntax. Operability accuracy – does it in fact work? Is this a functioning number? Is it a real address? And finally identity validations – is the person that’s identified or the entity that’s identified, is that the actual entity behind the domain?

So to date, that project has looked at syntax and operability accuracy, and they’ve actually issued several reports and I would commend you to the ICANN website on the WHOIS ARS system.

An interesting finding that we observed when we looked at the information that’s been collected is that new gTLDs actually showed higher syntax accuracy ratings for e-mail and telephone and lower syntax accuracy for postal address compared to legacy gTLDs. So that was one interesting category of information.

Some of the recommendations that we generated in our report regarding WHOIS is that we have this very large category of complaints that ICANN Compliance collects but we don’t have visibility into what particular aspects of WHOIS are being

complained about. So is it a complaint about bad information regarding syntax or is it a complaint that the identity is inaccurate? We don't know. And we think if there was that information available then the community could make decisions about, for example, whether the ARS project should continue on to the identity accuracy assessment phase. So that forms some of our recommendations in our report.

Next slide.

We also looked at safeguards in the sensitive regulated and highly regulated areas, and I'm going to talk about the sensitive and regulated strings first because those were subject to a certain set of safeguards. Highly regulated domains had an additional number of safeguards applicable to them. And this whole category was the subject of more safeguards than gTLDs that didn't fall within these buckets.

And this all arose out of very specific GAC advice that perceived that there were certain risks associated with strings linked to regulated or professional sectors. There was a sense that these gTLDs were very likely to invoke a level of implied trust from the public and therefore had higher levels of risk associated with consumer harm. So if consumers were going to be more likely to trust these gTLDs, they might be providing these gTLDs with

more sensitive information or engaging in more sensitive transactions, so these safeguards were aimed at that context.

Some of the safeguards that were applicable to the sensitive and regulated gTLDs required compliance with all applicable laws, and they're in a variety of areas – everything from privacy and data collection to organic farming and disclosure of sensitive data. And then there's a particular safeguard that really focuses on sensitive information as health and financial data. The safeguard has an explicit requirement to implement reasonable and appropriate security measures.

Next slide please.

Our Recommendation #23 really deals again with this topic that I previewed for you, which is, we have this whole trove of data in the complaints that are generated to ICANN Compliance but there's not sufficient transparency for the public to really figure out what is the precise subject matter of those complaints. And in this particular context, there's these requirements to comply with very specific laws but we don't know what law violation is being complained of and we don't know if the complaints actually even relate to one of the particular safeguards – and that's the protection of sensitive health or financial information. So if we're going to try as a review team to try and assess whether this safeguard has had an impact, we first need visibility

into the number of complaints that are received at the very least by ICANN Compliance regarding this topic. So our recommendation is really focused on getting more transparency in these complaints so that we can really begin to assess the effectiveness of this safeguard.

Next slide please.

For highly regulated strings, these were the types of gTLDs that are in the most highly regulated areas – charities, financial areas, for example – and these generated the most safeguards. There was a series of safeguards that required registrants in these areas to represent that they possessed the appropriate credentials and for a response about complaints, if there were complaints that said that the registrant in these domains didn't possess the appropriate credentials – for example, not a real accountant, not a real mortgage broker – and what we found is even though there was very explicit safeguards attempting to mitigate these risks, there's a lack of clarity about whether contracted parties are actually complying with these safeguards and whether the safeguards have been effective.

Next slide please.

So that generated a series of recommendations in our preliminary report to try and really get at whether these

safeguards are being complied with and if there are complaints regarding these particular areas because they may be the most harmful to the public.

So one of the things we recommended is to have an audit of a sampling of registrars and resellers offering highly regulated TLDs to see whether individuals without the proper credentials can actually buy one. That would be a good test. If the whole idea is to try and figure out whether this is being implemented, you might want to have someone who doesn't possess a credential try and buy a gTLD and see what happens.

Our other recommendations really relate to trying to get at the number of complaints again by seeking more detailed information about the subject matter of complaints to see if these particular gTLDs are generating complaints and then finally to look at abuse levels. There are gTLDs that have voluntarily decided to go beyond the contractual requirements and implement even more rigorous measures to proactively verify and validate credentials in these areas. The contractual requirement is that a registrant has to say, "I have all the necessary credentials," but there are a pool of registries operating these domains to have said, "We are going to take additional measures. We're going to actually verify and validate the credentials ourselves." It would be interesting to have

information comparing the abuse levels in those gTLDs, vis-à-vis gTLDs that don't engage in those additional steps.

Those are some of our recommendations.

Next slide. I think this is the question slide. Yes, this is the question slide.

This is just a sampling. We have a lot more safeguards that we assess and make recommendations about in the preliminary report. But this gives you a flavor of the type of approach we took and some of the recommendations we've generated.

If we don't have any questions from the forum then I'm going to pass the baton I believe over to my colleague Drew again – you're still not going to hear from David until a tiny bit later – to discuss the Public Interest Commitments.

DREW BAGLEY:

I'm now the opening act for David, though. It's quite an honor.

As Laureen mentioned, one of the safeguards that was developed as part of the New gTLD Program was the advent of Public Interest Commitments whereby a would-be registry operator, first in their application and then as a registry operator in their contract, could make binding commitments – binding because they're implemented into the contract – binding

commitments to do certain things that they would like to do, for example, with regard to DNS abuse, with regard to Rights Protection Mechanisms, with regard to registration restrictions, and whatnot. And so this was under the form of a very – and you can read more details in the draft report – but basically a very, very, hastily developed mechanism. And so because of the way that Public Interest Commitments came about, the registry applicants did not have much time. They had fewer than 30 days to submit Public Interest Commitments. And so we detailed that history and that definitely goes a bit into why we came to the conclusion about Public Interest Commitments turning out the way they did.

But what we ended up doing was we analyzed the contents of all of the voluntary Public Interest Commitments that were dealing with highly regulated TLDs, regulated TLDs, and the top 30 most popular of all TLDs for the new gTLDs. And from that, we determined that voluntary PICs absolutely vary in content and in the topic area they're covering, so as I mentioned, some operators would do something about DNS abuse, others would not. Others would focus on Rights Protection Mechanisms or other issues. And the purpose of the commitment was not always clear. Even if we were seeing that the commitment was to do something, it wasn't necessarily easy for us to figure out

how we would measure whether or not that was effective or how a community member would necessarily be able to determine whether the TLD was living up to that commitment they had made.

There also was no mechanism – in part because of what I described as the history of the voluntary PIC’s development – but there was no mechanism to ensure that a voluntary PIC was in fact doing something that was in the public interest in ensuring that it wasn’t a commitment to do something that was against the public interest.

Next slide please.

Therefore, from that we decided that first of all, in analyzing the Public Interest Commitments and therefore being able to scrutinize these better and ensure that these were a good thing or just determine what impact they were having, that they should be more accessible to the community. So all of the Public Interest Commitments are accessible by way of going to the actual contracts and viewing them, but as we learned – and even better than what as we learned as ICANN staff learned by helping us do this analysis – it was very difficult to do a comprehensive analysis and everything had to be manually extracted out of each of these agreements. So to even do a comparable analysis, when you were looking at one registry and

all of the registries that they operated, it was even difficult to do that comparative analysis to see if they had incorporated the same voluntary PICs amongst all of the new gTLDs that they operated and so on.

And so we believe that these voluntary PICs, because they are part of a safeguard, should be extracted from there and in a standalone database whereby community members could go and actually analyze them and access them quicker and that would help consumers trying to make trust judgements perhaps against various TLDs.

Also we recommend that future [gTLD] applicants should be required to state the goals of each of the PICs – just a very simple requirement just to state the intended goal and in the same way, put the PIC down assuming PICs are used in the future.

And then lastly, we believe that all voluntary PICs in the future should be submitted with ample time for various members of the community to be able to weigh in. And that would be effectively the safeguard of ensuring that they are in fact in the public interest as determined by the community.

We will have more analysis about PICs going forward before the final report. And as Laureen mentioned, one of the forms of analysis that we'll undertake is to compare those TLDs that had

voluntary PICs dealing with abuse with the actual abuse rates and see if there's any sort of correlation to learn there.

Now for the main attraction I will pass the baton to the David Taylor.

DAVID TAYLOR:

Thank you very much, Drew. I hope the anticipation hasn't killed anybody. It's pretty much killed me. I'm just grateful to have been left any time at all to speak. So there we go.

As you know, defensive registrations are one cost aspect we've looked at and Jordyn's covered those earlier. So the flipside of that is the protection of trademark rights holders and consumers by Rights Protection Mechanisms. We've been on a quest for data as well here to help assess the impact of the New gTLD Program both on cost and efforts and to protect trademarks.

How did we go about this? We've looked at the ICANN metrics from 2012 to 2015 which are based on the various Dispute Resolution Providers, WIPO statistics themselves in 2015, and we're updating both of those in the final report as well as looking at other sources. Talking of other sources, importantly we have the INTA Impact Study by Nielsen report. That questionnaire's been out and come back. The results are due on

the 3rd of April, so that anticipation is also killing me as well as everybody else.

The two most important findings – next slide if we can – really we’ve seen a rise in cases year on year [with] UDRP and URS combined, between 2% and 17% per year. That’s not necessarily surprising of course given that with new gTLDs the more domain names in the root. If we look at data for 2015 we saw that new gTLDs made up 6.5% of the total gTLD registrations, and I think one of the interesting initial findings is that UDRP disputes involving new gTLDs accounted for 10.5% of the cases at WIPO in particular in that same period. So we can tentatively conclude really that there’s proportionally more trademark infringement in new gTLDs compared to legacy TLDs. And looking at some of the figures we got for 2016, that seems to be indicating the same pattern.

Next slide please.

On the draft recommendations, as mentioned, we’re waiting for data from the INTA Impact Study in April. So that’s a key thing for us to formulate the final recommendations. We’re thinking it’s going to be so good we’re going to want to repeat it so that Recommendation #42 will be a study on the impact there going forward and repeated 18 months after the final report of the CCT Review Team and then to be repeating that regularly. And then

also full review of both the Trademark Clearinghouse and the URS, with the URS particularly the interoperability with the UDRP and this, of course, is being considered by the RPM PDP Working Group but there's obviously a need and we have this need for data to make recommendations.

And that's it. You've heard from me. I'll take any questions if there are any.

DENISE MICHEL:

David, are you also looking not only at the number of UDRP cases but the number of new gTLD domain names involved in the UDRP cases? In case the panelists aren't aware, there's been such high rates of abuse in some new gTLDs that we are much more regularly filing batch UDRP cases against some new gTLD registries. Thank you, David.

DAVID TAYLOR:

Thanks, Denise. And actually we could pop back to that last slide if we can, if we can go backwards as well as forwards. We are looking at that and on the pie chart I didn't specifically talk to it. I was going to – one more back – on the bottom right there there's the domain names involved in UDRP. We're looking at also as well in URSs. But you're looking at more domain names are involved in the new gTLDs' disputes than in the legacies'

TLDs. So that is interesting, and you've certainly picked up on that and we didn't plan that beforehand I hasten to point out for the –

JONATHAN ZUCK:

You want to advance the slides a few?

Finally, we were asked to look at the effectiveness of the application evaluation process. Obviously “effectiveness” is a fairly broad term, and so the review team decided to focus less on the inefficiencies of the process – which is really the purview of the Subsequent Procedures Working Group – and try to look at the inequities that might have occurred during the application evaluation process and suggest areas for improvement there.

One of the key statistics associated with the New gTLD Programs is the very limited participation in the program by players from the Global South. There were very few applicants that came from the Global South looking to have new strings, and so we wanted to delve into why that was the case and what it is that might be able to change that.

There were two things. One was an Applicant survey by Nielsen that asked the applicants questions about where things were difficult, whether they made use of outside consultants – which

they did quite a bit – etc., to try to see where some of these inequities might have arisen. And then, as I said, this meta survey or focus group of applicant cohorts or similar organizations [in] the Global South to the ones who did apply to ask them why they didn't and what it would take for them to apply in the future.

I would say that our number one recommendation is that the community needs to decide whether or not they want to make it a proactive objective to get applications from the Global South. And I think absent that commitment, it will really change the rest of the recommendations. And so we need to decide if that is in and of itself a worthy goal. It may instead be that getting more folks to register domains is a better goal in the Global South in the interim, or more people online, etc. That's a conversation the community needs to have.

But given an objective of getting more applicants, there's some things that we can do. One is a more comprehensive program of outreach. The outreach program was sort of cobbled together quickly and fairly late and didn't reach a lot of the right kinds of people and certainly didn't reach them in time for them to run things up the flagpole to raise the funds that would have been necessary for the program. More thought leader engagement

and more traditional media outreach seems imperative in the Global South.

Another thing that came up quite a bit in discussions with actors in the Global South is the need for case studies or business models for new gTLDs. Those that had heard of the program at all didn't have a sense of where the business model was. In fairness, there's plenty of people that participated in it that still don't know where the business model is, but I think we've gotten a little bit more knowledge and case studies of what seems to be working and not working, etc., what the minimum viable scale is. We're starting to learn more about that and we can share that information with potential applicants in a way that makes it less of a black hole for money in a part of the world to where the money's more precious. So building up these case studies seems to be important.

And then really revamping the support program. It was a very underutilized financial support program that involved a subsidy on the application fee, but there was also a program that I confess that I didn't even know existed, which was for sort of mentoring of applicants by members of the ICANN community and there was actually a list of people that signed up to be mentors. There was a list of people that signed up to get mentored. And somehow none of it happened. And so one of the

recommendations that we have is to have ICANN play a more direct coordination role so that the people that want help are actually introduced to the people that are offering to give it, etc. So that program had a number of people – 30 or so – people on each side of the fence there, many more than applied that were interested in that program, so we definitely want to see that program have a little more flesh on the bone.

Next slide please.

As we were looking at other areas in terms of dispute resolution and confusingly similar strings, there was a lot of inconsistency in these results and so our recommendations are about looking at how dispute resolution can be made to be more consistent, have the same groups evaluate the same kinds of disputes so that the answers are similar, we probably need to have an overall discussion in the Subsequent Procedures Working Group about the process including the potential of just prohibiting singulars and plurals because that seems to be a consensus in the community and yet somehow through a procedural process that happened during the initial round, that community consensus wasn't reflected in the outcome.

As far as GAC participation, it appears as though the early warning was largely successful even though there might have been some confusion about how much of a threat it

represented, for the most part, people were able to respond fairly specifically to GAC early warning, adjust their applications, either make a decision to withdraw when they could get most of their money back or make adjustments that allow their application to go forward. So there's, again, a rough consensus that the GAC early warning process was successful but advice itself appeared to be too late, too vague, etc., and so we've made a recommendation to the GAC to work on a template for advice, etc. to make it much more specific and actionable in the future going forward in future rounds.

Next slide.

Very quick because we're at the end, about some of the things we talked about – the application evaluation process. This is our timeline of next steps. We've just released a draft report. The public comment period is for the next 51 days. And then we'll be reviewing the public comments in early May. We should get the DNS Abuse Study and the INTA Survey in April and May and the DNS Abuse in June and hope to release a final report – fingers crossed – in July. Although the other thing that we're doing that's a little new is that some subset of the team will also stay on to help staff with implementation. Previous reviews there's been this notion of throwing the recommendations over the

transom and then complaining about implementation later. We're going to try to do something different this time around.

Next slide.

My recommendation to you is to read the report. And the rationale is that it may be long but if you get a cup of tea, it's really not that bad. We had someone read the report cover to cover, all 144 pages, and in theory we videotaped them doing it, so hopefully we'll be able to show this to you.

Okay. Needless to say, she was very enlightened and was able to gain great solace from her tea and the view out the window while reading it, so I recommend it to you. There's an Executive Summary that can lead you into the document and help you decide what parts of it are most interesting, and we really welcome your feedback on the recommendations – specificity, where we're too vague, etc. – so that we make this the best possible set of recommendations that it can be.

I appreciate you all coming out first thing on a Sunday morning as well. Hallelujah. Thank you.

[END OF TRANSCRIPTION]