ICANN Transcription
Transfer Policy Review PDP
Tuesday, 27 July 2021 at 16:00 UTC

Note: Although the transcription is largely accurate, in some cases it is incomplete or inaccurate due to inaudible passages or transcription errors. It is posted as an aid to understanding the proceedings at the meeting, but should not be treated as an authoritative record.
Attendance and recordings of the call are posted on agenda wiki page:
https://community.icann.org/x/YATpCQ
The recordings and transcriptions are posted on the GNSO Master Calendar Page: http://gnso.icann.org/en/group-activities/calendar

ANDREA GLANDON: Good morning, good afternoon, good evening. Welcome to the transfer policy review PDP working group call taking place on Tuesday the 27th of July 2021 at 16:00 UTC. In the interest of time, there'll be no roll call. Attendance will be taken by the Zoom room. If you were only on the telephone, could you please let yourselves be known now? Thank you. For today's call, we have apologies from Steinar Grøtterød. They have formally assigned Lutz Donnerhacke as their alternates for this call and for the remaining days of absence.

All members and alternates will be promoted to panelists. Members and any alternates who are replacing members when using the chat feature, please select panelists and attendees—or everyone if you have updated your Zoom—in order for all to see your chat. Observers will remain as attendees and will have access to view the chat only.

Alternates not replacing a member are not permitted to engage in the chat or use any of the other Zoom room functionalities such as

raising hands or agreeing and disagreeing. For today, we have three alternates who will be speaking and that has been approved.

If you are an alternate not replacing a member, please rename your line by adding three Zs before your name and add, in parentheses, alternate after your name which will drop your name to the bottom of the participant list.

To rename yourself in Zoom, hover over your name and click rename. As a reminder, an alternate assignment form must be formalized by way of the Google assignment form. The link is available in all meeting invite e-mails. Statements of interest must be kept up to date. if anyone has any updates to share, please raise your hand now or speak up.

If you need assistance updating your statements of interest, please e-mail the GNSO secretariat. Please remember to state your name before speaking for transcription. Recordings will be posted on the public Wiki space shortly after the end of the call. As a reminder, those who take part in the ICANN multi-stakeholder process are to comply with the expected standards of behavior. Thank you. Over to our chair, Roger Carney. Please begin.

ROGER CARNEY: Thank you. Welcome back, everyone, after a couple weeks off. Hopefully everybody is ready and energized to jump back into this. For a small group of people, actually, they've been working quite a bit over the last couple weeks. I think they had four or five meetings over the last few weeks just to discuss and talk through

the scenarios possible of auth code or AuthInfo changes. So I think that most of the day, we'll spend talking about what they've talked about and some of the high-level conclusions that they've come up with. But yeah, they'll present a little bit but mostly, time is allocated to discussion of all the points.

The group really had great discussions, and in a lot of these high-level points, there's obviously some additional work that needs to happen on them, and the small group recognized that and actually brought it up and discussed some of those. So I'm hoping the larger group will as well.

Before we jump into presentations, I just wanted to note that our project plan was introduced to GNSO Council. Greg did that for us, and in their last meeting, Berry actually introduced it to the whole group on the call and talked to it a bit. I don't remember if there were any questions. For some reason, no questions come to mind from Council on the project plan. So I just wanted everybody to know that now we're on the hook. So let's make sure we get our work done.

Okay, with all that behind us, I think we'll just go ahead and jump into some quick presentations. Jim Galvin is going to do some high-level discussion on principles and as you can see up on the screen, the last page or so of the working document has been updated with the small team's output. And then Jody will follow up with kind of just a strawman discussion on some more of the detailed items. So Jim, over to you.

JAMES GALVIN: Thanks, Roger. And as you said, many thanks to all of those who joined for the small team group work. We did have three or four meetings. I think I missed one of them. But we had a lot of really good discussions, and this is kind of where we got to for now. It's a good starting point. We spent a great deal of time talking about the way things are, the way things could be and what we're trying to achieve. So let's just walk through these. We'll see where we are.

I think this first bullet item is probably the most significant and most substantial and something that I want to call out, just speaking personally to folks to pay the most attention to. The others are more proposed details. We can certainly have a good deal more discussion about exactly what they look like.

I think this one really captures an important high-level overarching principle, and that is that the transfer authorization code—TAC—is an identity credential, and using it—and upon presentation, it represents that you are the owner of the corresponding domain name. So whatever domain name that you were issued that under, that's what that is.

And that's a really important principle. So I want to really highlight that for people to take in, really think about. I think that that's the way the AuthInfo code has been used for quite some time. But this really does pass into words the way that we've been working with it, and some of the discussions that we've had here as a group. So we tried to really capture that point, and I think it's essential that we get some agreement and thought about that point, because everything else kind of derives from it.

And yeah, now I'm seeing Farzaneh's comment in the chat room there about identity. Well, maybe we did [mess] on that one a little bit. That's a good point, we'll come back to that, I guess. The next one was, should only be at the registry when a transfer is in progress. What's going on here is two things.

One is we're trying to progress the principle that the TAC is essentially a one-time password and within that principle is the idea that there should be very little store, if any, of a TAC. And in particular, storing at the registry is probably one place that could actually be eliminated in the space. If you're looking to try and limit the amount of storage, that's a place which is easy to pull out of the mix. So that's sort of what behind this thought. It's been suggested that it isn't always what people want, there are registrars who had some other scenarios to work on, and so [we had to] give some thought to that.

The second reason why this is actually valuable is it also serves as a signaling mechanism. It makes it clear when transfers can occur or not occur. So if you always store it at the registry as part of the whole registration, then even the registry doesn't have any knowledge about it and it does create that vulnerability. If the TAC exists in general, then anybody could go anywhere and find a way to take it, and there's no way for the registry to manage that or control it. So that's kind of what's behind that particular principle. Just something to think about.

The next one down, must be unique per registrant and per domain name. This is really consistent with the rules that are there today, but making this very clear and explicit, one of the things that this does get in the way of that's worth thinking about is it does get in

the way of bulk transfers and bulk operations. So that's a due consideration here, something to think about. Will they be handled completely separately or do we really want to account for them here in some kind of different way? Should this rule be, this particular guidance element and principle here be tweaked a bit to allow for bulk transfers as opposed to bulk transfers being done completely differently? But that's really kind of what's behind there and that's the one competing thing that goes with it.

Next item down is, must not be retrievable from the registry. So this is just coming out of the idea that if the code is not going to be stored at the registry and if the code is going to be essentially a signaling mechanism that a transfer could be in progress, and also if you continue with the principle of a one-time password, then you want it to be not retrievable from the registry, which means you can set it at the registry. As a registrar, you would set it. And then the registry would be watching for it to come from some other registrar and then allow it to be used to conduct a transfer based on whatever else comes around from that.

But also, it means that at the registrar, if for some reason a registrant needed a new code or something else downstream needed a new code or they wanted to reconfirm the code, well, you just give them a new one. You don't actually try to give them the same code again. If there's a reason why it got lost or something like that, then you just generate a new one because you want to keep the uniqueness characteristic and you don't want to risk the idea that it could be picked up somewhere else and used. So it's there at the registry, it can be used if it's

presented, but as the registrar that set it, you can't get it back, you just would set a new one if that's what needs to happen.

The next principle was registrars should manage any TTL scheme. Again, this sort of aligns with the you don't store at the registry unless it's active. And now also since you can't retrieve it from the registry, you would need to set it to null if you don't want it to be used or you don't want the transfer to be active at that point in time.

But certainly, this leaves registrars—the important principle is that registrars remain completely in control of the transfer process. If you pass the TTL to the registry, then you're inviting the registry to be part of that role and process rather directly. They might just turn things off on you and that may or may not be what you expected. So that's just continuing that particular principle in this.

And then the last item here is it should be updated at the registry upon completion of owner change process. This is kind of open for discussion here. The issue here is that once the registry has completed the transfer, should the registry then set it to null so that it can't be transferred again? Or should we wait for the incumbent registrar to reset it? This seems like a good thing to do that the registry, once it's been used, again consistent with the one-time use principle, really want the registry to make sure that it can't be used again once it's completed that execution.

So that's our current draft set of principles to start with here as we consider this whole issue of use of the TAC, transfer authorization code. Happy to take any clarifying questions. And then otherwise,

I think, Roger, you have other plans for overall discussion, right? So, thanks.

ROGER CARNEY: Thanks, Jim. Really appreciate that. Yeah, and if there are any clarifying questions—I think once we see what Jody presents on strawman, it ties all ties all these kind of together to actually take that next step. But I think that then we'll lead into that discussion. I think the important part here was, like Jim said, that first bullet. I think you have to kind of buy into the fact that what we're trying to say is if you have this code, you're basically the owner. You can do what you need to do transfer-wise with it.

And some things came up with that as well as there was some discussion in the small group around the current AuthInfo being used for other things than transfers. And I think that that's some of the things that we want to pull out in discussion as well. Are there other things that are using AuthInfo? And do we need to look at those and see if they still are applicable, if they're using it maybe correctly or incorrectly, and take a look at the risk and benefit of assuming that the AuthInfo on the domain is for transfers only. And to this first bullet, if you do have that code, you can move it wherever you want to.

And again, I think some people mentioned in chat, it's not an identity, it just allows you to have that permission to move this wherever you want and realistically, make the registrant whoever you want it to be if you have this code.

And I think Sarah mentioned something about the creation of the transfer authorization code and possibly that being done by registrants themselves, which today, they're currently allowed to set their own or the registrar creates it and sets it. To my knowledge here, I think we always discussed this as the registrar doing this. We haven't introduced the fact of allowing the registrant to do it. Not that I don't think that they could or couldn't. I think that it is possible if we wanted that to happen, but our discussions to this point have been the registrar actually doing all that code creation.

All right, again, we'll continue this discussion once Jody presents the strawman, the more actionable items here. That may probably spur additional conversation. So I'll turn this over to Jody, and he can drive us through the strawman list. Thanks, Jody.

JODY KOLKER: Thanks, Roger. So the strawman list that we came up with is there. I'll just walk through these real quick. The first part of it is pretty obvious. The registrar would create the AuthInfo code and send it to the registry. We also thought hashing it at the registry would make it a little more secure as in no one should be able to break into the registry and grab auth codes in plaintext and start transferring domain names around.

The second one, I think we had a little more of back-and-forth on. Basically, the registrar should not store the AuthCode and it should be stored at the registry and only at the registry. I believe we had more of a problem with resellers of registrars in different

business models and how that domain code or the AuthInfo could be retrieved if it was not stored at the registrar.

We also had a problem—another issue that was brought up was how when a registrant has a problem trying to transfer a domain name, if the registrant or the registrar does not have the auth code, they can't prove that the registrant—they can't identify that the registrant has the right auth code. So it's useful to have that auth code stored at the registrar in order to debug customer issues. So I think that that's up for discussion here.

The third factor that we talked about was using the two-factor authentication to get the actual auth code. The registrant would be required to either have a cell phone or some other type of media, maybe WhatsApp or maybe a security question needs to be answered before the customer or the registrant is actually allowed to get the auth code.

A 32-character minimum was what we were hoping to get to, realizing though that some registries only have a 16-character minimum right now for the auth code.

The third one I think is very important also. We require different characters to be used: uppercase and lowercase letters. Currently right now, just letters are required by most registries, and then numbers and special characters. And then there are some registries that prohibit the use of dictionary words like password or peanut butter, anything like that.

And then one other consideration is homoglyphs. Basically, when you're talking to a registrant, trying to decide whether it's a zero or

a capital O, a one or a lowercase L or capital I, makes it difficult for the customer to be able to determine that. I know I've had this same problem on passwords or VIN numbers on cars, etc. So we'd like to say that you shouldn't be able to use those, in order to make it easier for registrants who are going to be typing them in if that's what they're doing instead of cutting and pasting from an e-mail or other type of social media that we have.

Then obviously, the registry would check to ensure the auth code meets—well, I shouldn't say obviously, but we'd like the registry to check to make sure the auth code meets the minimum requirements and deny any auth codes that are less than the minimum requirements.

This next item is a little tricky, timeout after a certain number of requests to initiate a transfer. Not sure timeout is the best phrase to use there, but basically, if there is a customer of a gaining registrar that is claiming to own a domain name and putting in several different AuthInfo codes or TACs in order to transfer the domain name that are basically trying to do a brute force of transferring the domain name, that the registry would stop those transfer requests for that domain name and possibly let the losing registrar or the current registrar know that that is happening on this domain name.

Then the last one, this was more of a question, I think, for debate, is, should the client transfer lock status and the server transfer lock status be removed before an auth code can be requested by a registrant? Basically, having another step in there to prove that the registrant actually owns the domain name. I think that's all I have, Roger, unless anyone has any questions. Thanks.

ROGER CARNEY: Thanks, Jody. I guess I don't have a whole lot to add. A lot of this was back-and-forth I would say on these as Jody mentioned, more of leaning one way or the other but still having open questions on most of these. And a lot of what I see in Chat is very similar, so it's good that we have this discussion within the whole group. But I did see Steve put his hand up. Steve.

STEVE CROCKER: Thank you very much. Jody, thanks for going through that list. As I listened to you go through that, it struck me that there are two competing factors—at least two that I'm going to talk about—under consideration.

One is, is the auth code strong enough? So that leads to the discussion of 16 characters versus 32 characters, choice of character sets and so forth.

The other consideration which is also very important is the usability. And part of that, you touched on with confusion of glyphs, of zero versus O and one versus L and so forth. But it also struck me that in pushing for very long strings, and then raising the prospect that humans are actually going to have to copy by hand, reading from one site, typing into another, that there is a very substantial usability issue there.

It's entirely reasonable to have both very strong authentication and a reasonably usable system, and I would just recommend that in designing this, both considerations be included. I don't have a specific formula to suggest, but just as an example, one could

imagine saying it consists of 16 characters, broken up into groups of four, displayed with spaces so that it's easy to see, and you restrict the character sets so that you don't have the glyph confusion. That would just be a quick reaction to illustrate the distinction between trying to get enough strength from a non-guessable point of view and a reasonably comfortable usability. Thank you.

ROGER CARNEY: Thanks, Steve. And to hit on the points you just mentioned, I think that some of that thought process was definitely in the small group, of trying to be secure and still make it usable. And I know that that discussions was in several of the tech ops in the past GDD summits, actually. They talked about that quite a bit, is making it so secure that it's almost not usable and then making it as usable as possible and then there's no security. So obviously, there's somewhere in the middle we can find that provides us with both of those. Thanks for that, Steve.

STEVE CROCKER: My reaction to 32 characters is, "Oh my god."

ROGER CARNEY: All right, any other comments? Volker, please.

VOLKER GREIMANN: Yes. I think one of the first questions we need to ask ourselves with regards to the provision of the auth code to the registrant is

how many entities in-between do we want to allow to have access to that auth code, bearing in mind that the business relationship between the registrar and the registrant is not necessarily a direct one?

So basically, we will need to decide as a group whether we want to have agents of registrants, be they law firms managing their domain names for their clients or resellers who are providing registrar services on behalf of the registrar, do we allow them to have that auth code or do we want direct provision of the auth code directly to the registrant?

I think this is an essential question that needs to be asked before we can look at all the other questions that are contained in this question whether the registrar has the storage, whether anybody else can store it. We first need to answer that essential elementary question, and once we have that, then we can move on and answer the other ones as well, because a lot of these answers might follow from that.

ROGER CARNEY: Thanks, Volker. Yeah, and I think when we looked at the principles, that first principle of if we all accept the fact that if you have this code, you can transfer it—and that's all that's really said, if you get it then you can transfer it. And if everybody accepts that, I think that you're right, there seems to be enough evidence and [use] that saying storage can't be done is probably not a valid concept. I think the limited storage note just on who stores it but how long becomes more important than saying, "Yeah, don't store it."

Obviously, the fewer times you store it somewhere, the better security-wise, but I think that there'll be times when that has to be at least [inaudible] from one area to another, so you are technically storing it. So I think that to your point, Volker, yeah, we have to get to that spot where—I don't think we can say no storage, it's just how we want that to happen, how long do we put good parameters around, "Okay, only keep this for as long as needed, only keep it for X number of days" and things like that. Thanks, Volker.

KEIRON TOBIN: Thank you. I'm just trying to wrap my head around in terms of maybe I'm missing something here, but if there was a two-step authentication and the registrar create the auth code, if the registrant was to lose their kind of two-step authentication or 16-digit code in regards to that, how long would it take to reset? What are we looking at in terms of, would we need to reach out to the registry? It creates a kind of issue in terms of customer experience. And have we looked into any potential in regards to that?

ROGER CARNEY: Thanks, Keiron. Yeah, and I think that when we talk about the two-step, we have to think about where that's appropriate. I think today, if you look at most of it, the two-step is into the registrar panel, whatever you want to call it, their domain management panel. That's where the two-step is. But once they have the auth code itself or the ... There is no two-step after that. Once they have it, they're done. And I think that from what I've understood

anyway up ‘til now, that’s still the idea. There's no two-step after achieving or getting the auth code. Others can jump in if they think that that’s a good idea or not, but to me, that two-factor was always before the auth code is generated and given.

Lutz, please go ahead.

LUTZ DONNERHACKE: I have some questions. The most important point here is the discussion about the ecosystem outside of the registrar and registry. The whole ecosystem [inaudible] the registrants, because they are not covered by the contracts of ICANN, they are living in their own world, they are unable to do a two-factor authentication simply because they're technically illiterate in most cases if you do not talk about large registrants but the small ones, the millions of having one or two domain names at all, and they are all using resellers, they do not go to a registrar itself. They use a chain of resellers.

So there are a lot of questions arising here. First is about storage of AuthInfo code. [inaudible] in order to gain the AuthInfo code from the registrar to the registrant and in most cases, they are sent by e-mail, they are stored in various ticket systems and might be retrieved many times by many people. That’s reality, so we have to deal with it. And the second point here is that if we’re going to make technical difficulties in order to access the codes or to transfer the codes, we will lose a lot of such people. Please keep this in mind. Thank you very much.

ROGER CARNEY: Thanks, Lutz. And I agree, the reality is we won't be able to say "Don't store this." It's coming up with the good practice around storing and keeping that safe. It's not safe for anybody except for you're trying to make the registrant, their value or their ownership in that domain safe. Okay, any other comments, questions?

THEO GEURTS: Hey Roger.

ROGER CARNEY: Hey Theo.

THEO GEURTS: I'm sorry, I can't raise my hand for some reason. Great work from the small team. I do have a couple of questions though, or maybe some clarifications. I really like the way where it was heading, and giving the registrant the auth codes and being able to transfer it all. And as mentioned, bulk transfers would be an issue, but if this group can work on alternative solutions when it comes to bulk transfers, I would be totally fine with it to keep going where we are heading now, but also in addition, to do some extra work to get a working bulk transfer system in place so that resellers are not being affected negatively about this. So good work on that as a group. Fine, great.

Jody explained something regarding multifactor authentication, being able to unlock the domain name. Not sure if I'm paraphrasing that right since I was looking for the raise hand button through the presentation. But my issue with it is I think a lot

of us already offer MFA, but usually, the user uptake who actually turn it on is very low. So it could be a very great way to have multifactor authentication really get adopted if it becomes a requirement. It could also lead to an uptick in support. So maybe a little bit of thinking there in that area is also something.

And there was also mention of the transfer lock needs to be removed prior to transfer. I didn't understand that very clear. Oh, yeah, the updating and the requesting for the auth code. What we see is that people, our resellers or the customers of our resellers turn that on and off for no obvious reason. And I'm wondering if that is really great way what we are looking at now.

What we see is some resellers turn it off by default, they don't want to have the locks on it for whatever reason. It could be they are also an Italian registrar and then you can't have those locks on those domain names anyways so they don't put them on gTLDs also.

So there's a lot of use cases there regarding when a transfer lock is present or not. So I'm just putting that out there. Thanks.

ROGER CARNEY: Thanks, Theo. And I'll let any of the small team speak to that specifically, but my thought was when we were talking about that was it's more of a timing issue and should you be able to get, ask, request and actually receive or not an auth code when a lock is on it, should the lock have to come off to get that and things like that. Sorry, Sarah, I'm going to jump you and see if Jothan has a response to that.

JOTHAN FRAKES: Sarah, did you have something? Just so I don't have the queue on you here. So I had suggested this, and clearly, as you look through the strawman list, there are things that may be in conflict with each other. So for example, the second bullet point saying that the auth code might only be stored at the registry.

So if that were the case, the transfer lock status being on might affect if you can request the auth code from the registry was that suggestion. Clearly, there's a diversity to how transfer lock is used currently. But if that transfer lock is set, would you want to be able to use that as some way to restrict whether or not the registrar could be setting the AuthInfo code?

So those were some of the things to just consider in the strawman list. And hopefully, that helps with the conversation. Thank you.

ROGER CARNEY: Thanks, Jothan. Sarah, please go ahead.

SARAH WYLD: Hi. Good afternoon. Hope you can hear me. I wanted to talk a little bit about actually the second principle, which is a little bit higher up in the document and how it relates to the second point on the strawman list, and so apologies to those who have already read my comment, but this is what I thought when I first said this earlier today.

So, how exactly would the registry know when a transfer is in progress? Or how does the registrar know when to send the auth code to the registry, the losing registrar? So as I wrote down, the domain owner unlocks their domain name. Is that the trigger to send the code to the registry? Because just unlocking it doesn't mean necessarily that they're going to transfer it later.

Also, we should keep in mind that some registrars bundle together the different EPP status locks on a domain. So there's client transfer prohibited. There's also client update prohibited. So we often see a domain owner unlock their domain to remove both of those locks together, change their nameservers and then lock it again. They don't want to do a transfer, they just need to unlock the domain to make the other update. Right?

So, is unlocking it going to be the trigger to send it to the registry? And if not, then at what point would they do it? Because they can't wait for the domain owner to take that code and give it to the new registrar, because then the new registrar has nothing to verify it against in the registry system. It has to already be there.

So I'm not super clear on exactly how that would work in real life. And also, it doesn't quite seem to align with the second thing on the strawman list that says—and so maybe I misunderstood what that point is. I added in the words "it is stored at the registry" just to balance out it's not stored at the registrar. I think that's what the intent was, but that says it is there, and then up above we say the registry only has it when the transfers are in progress.

So I feel like maybe I'm missing something or maybe it just needs discussion. Thank you very much.

ROGER CARNEY: Thanks, Sarah. And I'll let any of the small team members talk too, but I was just going to add in that I think the idea is when the registrar provides the registrant the new TAC, just prior to that is when they stored it at the registry. And then they'll manage that however they see fit based on other rules, but that's when the registry would ... I don't know. They're not going to watch them. Whatever. However they know. But that's when the ecosystem would know that a transfer request has been put in. And again, it's not that a transfer's been truly requested, it's just that the transfer request has started.

And I think that that was the timing, was when the TAC is provided to the registrant, that's when it's written to the registry. And as many people have noted, I'm sure that that's going to get stored at various spots. Maybe customer service has it for any issues of readability or whatever and to help the registrant or reseller to reseller, reseller to registrar, however that works, may need to do that as well.

Again, I would propose this language around keeping it as minimum as possible but, again, the group can make those decisions.

SARAH WYLD: Thanks, Roger. Can I speak to that again?

ROGER CARNEY: You bet.

SARAH WYLD: So I think in a lot of different provider control panels, but registrar and reseller, right now the way it works is that the auth code is always there, always visible. So they don't know when the domain owner retrieves that code. So what we would be doing would be updating to a system where there is no code set until the domain owner goes into their control panel and does something to generate one.

And if that's the case, I feel like that's just not clear here. But it would be quite a change. Thank you.

ROGER CARNEY: Thanks, Sarah. And I think that so far through the discussions, that was the idea, was only at the time of request would an auth code or tac actually be available, created. Then again, I'll leave that—any small team members want to talk about it too? Jim, please go ahead.

JAMES GALVIN: Thanks, Roger. I understand Sarah's concern, and it occurs to me that I don't think we made clear one particular point in this discussion. This set of principles and the strawman—there are obviously some details which makes them look like they're out of sync with each other. I guess those of us who were writing these things were a little too close to it. So Sarah is highlighting a bit of confusing points in all of that.

But I think the larger point to keep in mind is yes, this is a proposal, and this proposal does not necessarily match 100% what everybody's doing, which means part of what we need to do here, Roger's already called out the question of what do people use the AuthInfo code for other than transfers.

We've heard anecdotally that there are some other uses, and it's important to take all of that in and understand it and what changes people might be looking for if they're using it or if some of our principles here are going to impact that. We need to understand that.

But the other key thing here is this does not necessarily represent today's state of the art. There's a certain expectation that some set of people—maybe most, maybe not, I don't know, that's open for question—will probably have to change something. Hopefully just a little something, not a lot of something, in order to align with these principles if these principles are what the group adopts.

So one of the overarching questions here for the entire group is, do these principles conflict with where you are in a way that we'd want you to seek a different solution, that we have to adjust these principles so that there isn't a fundamental change for you. Hopefully, they don't represent a fundamental change, but that's what we have to call out here.

And I think Sarah was just about going down the path of, "Wait a minute. I think one of these things really gets in the way of our standard business processes." And it's useful to call that out explicitly, get it on the table so that we can think about how we

might tweak these things to not make such a critical change to how you do things.

So I hope that helps the discussion. I wanted to build on what Sarah was commenting on and call that out for everyone. Thanks.

ROGER CARNEY: Thanks, Jim. And to your point and to Sarah's point, I don't want to be disruptive to be disruptive. If we're going to be disruptive, we want it to be purposeful and only make those changes that are necessary to improve this process to the way we think it should be. Kristian, please go ahead.

KRISTIAN ØRMEN: Thank you. I really think it would improve the security if we were at a point where the auth ID would only exist when we created it and for a limited time, so we go away from the current practice where auth ID is always there.

So I think that would improve the security of the auth ID and also be in line with this, that it should be like a one-time password. But then we should really have a time to live on it. And I really think the security around that would be better by being handled in the registry system and not to say that it should be handled by the registrars.

And one comment about what auth ID is also used for, and that would be transfers between resellers. Quite often, one reseller get in a transfer, maybe that domain is already at the registrar they're

using, so that auth ID would then be just used for a transfer between different resellers at the same registrar. Thank you.

ROGER CARNEY: Thanks, Kristian. And again, going down what Jim and Sarah were mentioning, was how destructive this may or may not be, one thing to keep in mind is we're talking about the TAC here, the AuthInfo in kind of in isolation, but this discussion will lead us into the discussion on losing and gaining FOAs as well and all those things will come together to make one solid recommendation set. So I think again, we're focused here on the TAC, but obviously, things that we're doing here are going to affect those other things. So we just need to keep that in the back of our minds, not in the front of our mind. Yes, Jim, not destructive. Thank you. That is our goal.

Okay, again, as Jim mentioned—and I think Sarah kind of found a few sub-bullets or additional bullets here. I think that was the purpose, really, is here's what the thoughts are, and what's missing, what needs further discussion, can we accept the bullet one that says, hey—without the identity in there, but can we accept bullet one saying if someone has this code, they can transfer it and that's just a fact that they're showing—and again, whatever the previous registrant was, no one will know. The gaining registrar doesn't need to know. The registrant itself could change theoretically. There's no way to balance or to figure that part out.

So I think that we need to look at all these items and say, okay, what's missing in between? And to what Kristian just mentioned,

let's add on to okay, and it gets used for this. From reseller to reseller, that's important to note. If that auth code only exists during a transfer, then, say, it's going from reseller to reseller, then the reseller has to request it from the registrar, registrar's got to set it and provide it back to the reseller that can provide it to the reseller to move it. So again, thought process on how disruptive that may be. We need to add in all those use cases and look at each one and say, "Okay, can we solve that or not?" Jothan, please.

JOTHAN FRAKES: Thank you. Just a quick comment is that some of the things—I think if we had a little bit more time, we may have put some swim lanes that would help to clearly identify what is going on at the incumbent registrar, what's going on at the registry, what's going on at the gaining registrar, and some of the principles around maybe things that were the user interface at the registrar or requirements at the registrar as opposed to things going on inside the SRS.

I think you see a couple different times here where we talk about one-time password or two-factor authentication, those are things that would be really the user interface at the registrar, either gaining or losing, in order to help ensure that there are layers of security and protections for that information.

But in other cases, we're talking about maybe what the minimum complexity or length of an auth code is, and that might be something that is interactive. Maybe it's set as business logic at the registry and the set command might fail if it doesn't have a

certain level of uniqueness. We didn't want to necessarily be rote about this, and I think if we had more time, we might have made this a bit more clear.

But high-level, even though as Jim points out that these may not be 100% aligned, we tried to capture all the different ideas. And as I mentioned in the chat, I think our objective as a small team was really we gather up the LEGOs, make sure it's only LEGOs and then dump those out onto the table and let the working group do their thing. Thank you.

ROGER CARNEY: Thanks, Jothan. That's a perfect way to put it. Berry, please go ahead.

BERRY COBB: Thank you, Roger. And for the record, it's LEGO even in plural, not LEGOs. But anyway. Jothan, I appreciate your point about the swim lanes. And staff did some digging in the background about what current documentation or diagrams do we have about the transfer process. And unfortunately, we've seen very high-level material, a lot of material that's probably more registrant-facing than what would be useful for our purposes in discussions here. And I even went back into the archives back to IRTP part B where I participated on that long before my current role, and I actually put together a swim lane process flow diagram that had all of the roles, registrant, gaining, losing, registrar, registry. There was a swim lane for the systems or for EPP.

And if the group thinks that that's useful, I can try to knock the dust off of that and try to bring it more up to date, because it's definitely—it wasn't correct at the time and now ten years have g one by, so it's even more incorrect. But try to develop a working draft that we can use to facilitate the discussions about what we're dealing with with current state today. But I think what would be most useful is after we've started to formulate preliminary recommendations about what's going to change or what's going to stay the same, that that become a work product that we deliver with the final report that can be used not only for education and implementation but maybe to help facilitate future policy discussions if they're ever needed. Thanks.

ROGER CARNEY: Thanks, Berry. Theo, please go ahead.

THEO GEURTS: I'm not sure if going back to all these use cases is a very good way to spend our time. I suspect that if you would go really deep on all these use cases, you will end up with use case after use case after use case and I'm still pretty sure that we won't cover them all because there's always regions of the world which are not covered by this working group where things are going very differently, usually. And then you create a lot of work.

I would just propose a different path, and that is make sure that we set our goals very—put out our goals, see what they are and make sure that our language around those goals, what we try to achieve, what the purposes are and then make sure they are

somewhat technology agnostic and somewhat policy agnostic so that there is enough room for all the registrars to achieve the goals that we want to set as a working group but not restrict registrars or any other contracted parties with a whole set of very specific policy requirements which will limit several use cases or will put registrars in an impossible position due to very overly specific use of the policy requirements.

So I would keep it very open and make sure that the language that we use—just make sure that we hit our goals as a working group and that registrars can implement it for everyone. Thanks.

ROGER CARNEY: Thanks, Theo. And as Berry described, I was thinking, yeah, it seems like a perfect document for a reactive kind of—and as Berry mentioned, maybe it's more of an educational implementation kind of thing later on, not necessarily driving this group but helping understand what our recommendations kind of led to, kind of idea. Any other comments, questions?

And Lutz had a good point made in chat, if there's a fallback procedure for a failed or wrong transfer. I think that that's something we still need to discuss later on. I think that obviously, that still has to happen, have a good way to recover any issues of a transfer. Maybe it was illegal or just wasn't supposed to happen. But I think we have to cover that topic as well. Jothan, please go ahead.

JOTHAN FRAKES: Thank you. Yeah, and one thing that came up in the small group's discussion had to do with how certain aspects of what we're talking about just in the scope of the AuthInfo code is that there's pieces as we go through this whole working group that are going to be interdependent.

One other element of the working group discussion will have to do with partial bulk or bulk-type transfers. And one of the situations or circumstances of use of auth codes was that there are real-world scenarios where someone may be consolidating their domain names from across a variety, a diversity of different registrars, could happen from company acquisition or a variety of different reasons, but some poor person is tasked with I want all these domains to end up at registrar X. And that becomes a process where that person who has that responsibility will have to go and kind of accumulate all of these different auth codes in some manner and put them into effect in some form of transfer.

And so as we look at this, we do want to bear in mind that we may be having conversations elsewhere about partial bulk or things where maybe that use case is or isn't covered. So there are pieces here where for example in the scenario I described, somebody may need to gather, accumulate, store and then process all at once and track those transfers as they come from various sources.

I hope I articulated that well. I'm low on my caffeine here. Thank you.

ROGER CARNEY: Thanks, Jothan. I think we're doing what I would hope we would do, kind of looking at the principles of the strawman list as, as Jothan put it, several LEGOs here—LEGO. And the small group knew that there'd be some things missing or even questions about the things that are there. And that was the whole purpose, is to put these down.

I've always found that people find it much harder to write than to edit. It seems like everybody is a great editor but not necessarily a great writer. So I want to thank the small team for putting this down so that we can start to edit this, throw darts at it, see where the holes are, uncover any issues and start documenting what those are so that we can evaluate each one of those and see what the impact is.

So again, I think that we're doing that on this call well, and my goal is between now and next call, one of the big homeworks will be looking at these two things, the principles and the strawmen, and everyone put in there those ideas and concerns.

Steve put in that 16 seems more reasonable. Maybe in quads spaced out evenly. However you want to do that. And others think that—I think it was Kristian that mentioned that the reseller may be sharing of the TAC. So put that in there so that we know that and we can look at that and see how that's impacted on any decision. So again, we're working through that now and staff is documenting some of those things, but I want everybody to take a look at these two lists, the principles and strawmen and do those things over the next week as well so that we have a lot of good comments in there and a lot of good direction.

So not trying to close the call here, I just wanted to make that known so that everybody is aware of where we're heading. So, any other comments, questions here?

Okay, so I haven't heard anybody disagree with the original premise, and I think that that's kind of important if we got that down and no one disagrees with the fact that if you have the TAC, then you are considered—I don't know if you're considered the owner or not, but the simple fact is if you have the TAC, you can move that domain wherever you want and in that process, possibly change the registrant as well. So updated wording here. Thank you.

Okay, so I like that everybody agrees to that first one. The second one, it sounded like there was maybe—not necessarily disagreement, but more continued discussion for understanding—I know Sarah had put in a good comment here about what that is. And I'll open this up. That was my understanding, was the goal from this group was when a TAC is created, when it's supplied to the registrant, that's the only time that it lives at the registry and then after a certain period of time, it'll actually disappear from the registry as well.

We haven't really got into the TTL discussions yet, so something we can, I guess, [inaudible]. In the small team, there were some discussions about TTL and who would manage the TTL, and it sounded like that there may be multiple levels of TTL. Maybe there's a maximum system TTL when it's created, that the registrars or the registries are sort of responsible for. So when a TAC is written to the registry, it only lives for 15 days. And I'm throwing 15 days in as the X value, that once the 15 days is gone

and that code is there, if it's tried to be used by a gaining registrar, it'll fail because it's too old.

Additionally, I think there was discussions of registrars managing the TTL at a finer grain level, not as a maximum but as maybe more of a flexible value that for maybe high-value domains, it's only two days or something and that the registrar would just go blank the TAC at the registry when that actually expired so the registry didn't have to manage that. I don't know, thoughts from people? Does it make sense to have two levels? Does it make sense that it's identified as being managed by two different people? The registry managing the maximum, the registrar managing something inside that. Kristian, please go ahead.

KRISTIAN ØRMEN: I can at least live with that solution as long as there's a max TTL on the registry level. I would be even more happy if I could also just set the TTL where when I create the auth ID at the registry level, but I could live with the solution that you just said.

ROGER CARNEY: Thanks, Kristian. Tom, please go ahead.

THOMAS KELLER: Thank you. I could live with it, but I would pretty much prefer if it's done by the registry just for matter of transparency and consistency. We have fewer and fewer registry operators and we have more and more registrars. So my big fear is actually that you will have, as always with policy, different kinds of understanding

[across registrars,] how to treat that, and there will be compliance case after compliance case to make sure that all comply with the same rules even if they're written in text.

So I do understand that registry find it hard to actually implement something in their systems, but just from economical tradeoff, I think it's much more preferable in the overall cost scenario that the fewer parties are doing it instead of the many. Thank you.

ROGER CARNEY: Thanks, Tom. And I think that's where you have to start looking at, to your point, the consistency or it's not a big, it's a feature kind of flexibility into it. It can happen today as well, but I'm also thinking of scenarios where the registrant may think that their code got compromised so they call the registrar back and say, "Okay, let's stop this or give me a new one," and that process starts over. Jothan, please go ahead.

JOTHAN FRAKES: Thank you. So for what I found in going across a variety of different registrar interfaces is that some will set the auth code immediately upon request or only share it. Others would let you set it and it's just set. And it's the transfer lock that is kind of guarding the name. I've seen that implemented in different ways across different registrars. So that's kind of the status quo we're working with here. Just wanted to make that on the record.

The other was a quick comment that there are some scenarios that might make it legitimate to have a very brief time to live as a minimum if the registrars are setting it, but I could think of

scenarios where a registrar might set that perhaps aggressively low to make it incredibly difficult for a registrant to move to a new registrar. So we may want to define some reasonable minimum that must be there. Thank you.

ROGER CARNEY: Thanks, Jothan, and that's a great point. I think you're right. I think an easy path for a bad acting registrar would be to set it to a minute, ten minutes even, something that's not reasonable, to go to other sites and finish the transaction at. So I think yeah, there would have to be some kind of minimum put in place. Jim, please go ahead.

JAMES GALVIN: Thanks, Roger. On this issue of should there be a TTL stored at the registry and should the registry enforce it, I'd like to change the discussion and reframe the question. It was suggested on the one hand that it would be easier for the registries to implement this because there's so much fewer of them than there are registrars. So there's a greater incidence of inconsistent implementation.

And while all that could be true, and whether or not it's harder at the registry I don't think is really the issue that we really should be dealing with here, what I offer as a question is, what is it we're trying to achieve here? What is the goal that we're trying to achieve? Who benefits from having something in one place or another? And really, we're talking about benefits that [inaudible] to the registrant more than anything else.

So, what is the security feature, if you will, that we get with one mechanism or another? So if you store the TTL at the registry an then the registry is going to enforce it and so the registry will just turn it off when whatever we decide that [inaudible] is, well, that means that the registrar loses control of any kind of options that they might want to have in terms of giving longer or shorter TTL periods for different customers.

The other thing that happens, though, is the registry will just turn it off and you have to track that as a registrar regardless, because now what happens is they're just going to turn it off, the registry is not going to have the relationship with the registrant. If they were slow in doing whatever they were trying to do, and in going to new gaining registrar, you're still going to get that phone call. You're still going to get that particular customer service issue.

And they're going to reach out to you and contact you in the same way that, you're right, registrants have to learn the proper way to deal with the system. Yes, registrars are going to have to teach registrants or make available to them the right information.

I think that there are failure points in all of these modes. So I would focus less on where the failure points are because I think those compliance issues exist regardless. Customer service issues exist regardless, you're just shifting the from one form to another. I really do believe that in the grand scheme of things.

I tend to focus on what is the benefit in terms of security that the registrant gets when you do it, what do we consider to be the best guidance in terms of security practices and the best guidance in how do we achieve that? And then we really just need to move

ourselves in that direction. And obviously, it's not going to be perfect for some time, but hopefully, we want to get to a better place and we can all achieve that.

ICANN Compliance should be left to do their job, whatever we decide that needs to be, as we focus here on what's best. Thanks.

ROGER CARNEY: Thanks, Jim. Tom, please go ahead.

THOMAS KELLER: Thank you. This seems to be a bit of a hot potato, so I think we've been there before [at a point you say] maybe it helps that we formulate the policy first, so what should be the TTL about, how long should it be, what are the maximum set, the minimum set? How short is it? How the exact flow we think there should be. And yes, there will always be registrars that do it a bit differently, but I think there should be a certain expectation set how the process looks like, because what we're trying is to improve the usability for the end customer.

So, why don't we just do that first? And once we have that process, we can still decide who's going to implement it.

ROGER CARNEY: Thanks, Tom. Kristian, please go ahead.

KRISTIAN ØRMEN: Thank you. So I think the most important to start with is to agree if or if not we should have a TTL on auth IDs. I think we should have one. If we have a TTL, the way the transfer I think is going to work is that the registrant will give the gaining registrar the auth ID, the auth ID will then send the transfer to the registry to get that auth ID validated. The registry in my mind can only validate that auth ID if they know if the auth ID has expired or not. and because of that, it's why the date of that auth ID needs to be in the registry.

ROGER CARNEY: Thanks, Kristian. And I think the point you made and the point Tom has made is, yeah, let's decide if we think TTL is a needed feature. And from our discussion so far, it sounds like the majority of people think it's a good feature to have and there probably should be an upper and a lower limit, but then how the management is still a bigger discussion, but if we agree that TTL is something we should do, set those bounds and then we can get into those later discussions. Thanks. Volker, please go ahead.

VOLKER GREIMANN: Yes. Just to reiterate a little point from the small team discussions, which is if we set a TTL, we always place some part of the risk with the registrant. Not the risk of the auth code getting known or abused but rather, the risk of situations like the slow motion failure of Net4India where ICANN took about a year to deaccredit them and get the transfer process installed. In the meantime, registrants were unable to transfer their domain names because the systems had been turned off.

Had they had the ability to query the auth code and store that auth code in a secure facility like a password manager that they have on their PC anyway, they would have had access to that auth code. And now if we propose that an auth code that they might have requested a year ago or two years ago, just to be safe against certain eventualities, suddenly becomes no longer valid, their registrar or reseller goes belly up and they don't know where to turn to ... We just need to be certain that we do not exchange one risk with another risk: one risk which is the theoretical risk of the auth code getting read out by the registry or the registrar, by the third party that has access to their systems, or the very real risk of registrar failure, of reseller failure, of any failure in-between where the registrant ends up holding the short end of the stick. And I think that's also something that we need to prevent. Thank you.

ROGER CARNEY: Thanks, Volker. Jody, please go ahead.

JODY KOLKER: Thanks, Roger. So I've heard about this before, and it's not just the auth code that you need in order for the registrant to transfer. The domain has to be unlocked at the registrar. And if the registrar's failing and you can't get a hold of them, just having that auth code won't allow you to transfer that domain name. It also has to be unlocked, and if the registrar has failed and is no longer active, and the domain's locked, that auth code won't do any good for anyone. You can't unlock the domain name at the registry, and

I don't think anyone would call the registry and say, "Please unlock my domain name" without the registrar being involved.

So I think that we need to make sure that the auth code is just not keys to the kingdom. There's one other thing that has to be done too. In a registrar failure, if the registrar has all the domains locked, there's no way that domain is going to be transferred even if you have the auth code.

ROGER CARNEY: Thanks, Jody. A good point. And maybe it's a good point against our number one principle here of if you have the auth code, then you can move it, which isn't actually true from what you said. But I think it goes to what Jothan actually brought up as the last strawman idea, I think, was the timing of locks versus auth codes and if a registrar does provide the auth code to the registrant, then that doesn't assume ... it makes sure that the locks are off so that it can be transferred. And again, that would be some policy that we would have to make that, okay, once an auth code is set at the registry, the locks are off. It would have to be part of that, but to your point, Jody, the auth code is not the only key to it. So I think we have to be careful there.

Okay. Any other comments, discussion on these items? I think that second to last bullet on principles, on TTL, registrars should manage this, I think what the group has said is we should have TTLs—and I'm just going to throw this out there, I think that the group is saying we should have a maximum and a minimum TTL as well. But how that TTL is managed, the group has not agreed

to yet. But the group has agreed that TTL should exist and there should be a minimum and a maximum. Kristian, please go ahead.

KRISTIAN ØRMEN: When you're talking about minimum, I wanted to be sure that we had somewhere added that we should always be allowed to override it with an auth code, and I'm not sure if that is in the document. And if not, maybe it should be added, because if you add a minimum TTL on an auth ID and it somehow gets compromised, we need to be able to invalidate it no matter the minimum.

ROGER CARNEY: Yeah, I agree. Thanks, Kristian. Okay. And I think that, again, the last bullet of the strawman list that Jothan threw on there, I think in the last meeting actually, I think is coming out to play more and more as we talk through the whole process and again how the locks work with the timing of an auth code and things of that nature. So I think we have to think about that last bullet a little more and how that would work out.

Okay, I don't know if we had any other solid changes to this. But again, I think the homework from now until next week is to put your comments in here. If something needs to be added, something was missed, if you think something should be in here, let's put it in here. And again, if there's any confusion, Sarah did a great job noting on the second point and getting that clarified so that we understand what we're saying and then we can make those decisions of, okay, does it break something else?

So I think that's the homework. I think we have ten minutes to go. I don't know if anyone has anything else they want to say or not. Otherwise, I'll turn this back to staff to see if they have any closing. Staff, anything that we need?

EMILY BARABAS: Hi Roger. Nothing I can think of, unless others want to step in. I guess we'll follow up with notes and actions on homework, and I think Andrea can close the call. Thanks.

ROGER CARNEY: Great. Thanks, Emily.

ANDREA GLANDON: Thank you. This concludes today's conference. Please remember to disconnect all lines and have a wonderful rest of your day.

ROGER CARNEY: Thanks, everybody.

[END OF TRANSCRIPT]