MONTREAL – GNSO - CPH TechOps Meeting
Sunday, November 3, 2019 – 09:00 to 10:15 EDT
ICANN66 | Montréal, Canada

MARC ANDERSON: While we're waiting to get started, I'll just say there's plenty of seats at the table if anybody wants to come up. I encourage you to sit at the table. This means you, Quoc.

ZOE BONYTHON: Okay, we're going to get started now, if we can start the recording, please.

MARC ANDERSON: Good morning, everybody. My name is Marc Anderson and this is Tobias. We're the co-chairs of the CPH TechOps group. So, hopefully everybody is in the right room. Welcome to ICANN 66 and the meeting of the TechOps group. What I'd like to do, to start things off, is get everybody to introduce themselves real quick. We're a good-sized group of people in here, and there's a lot of faces I recognize, but there's some faces I don't recognize. So, if I can pick on everybody to introduce yourself. Let us know where you're from. I'll start down there at that end of the room.

[BENET NORDRIG]: [Benet Nordrig], coming from Sweden, Norway, Denmark—Scandinavia.

[RIGA POVE]: Hi. I'm [Riga Pove]. I'm from one.com.

KRISTIAN ØRMEN: My name is Kristian Ørmen. I'm from Larsen Data in Denmark.

GREG DIBIASE: Greg DiBiase, Amazon Registrar.

RASAM HAFEZI: Rasam Hafezi, Amazon Registrar.

VLAD DINCULESCU: Vlad Dinculescu, DNS Africa.

OWEN SMIGELSKI: Owen Smigelski, Namecheap.

BEN MCILWAIN: Ben McIlwain, Google Registry.

ZOE BONYTHON: I'm Zoe Bonython. I'm the Registrar Secretariat, but I do do the support for TechOps.

SARAH WYLD: Hi, I'm Sarah Wyld. I'm from Tucows.

ANTHONY EDEN: Anthony Eden from DNSimple.

MARC ANDERSON: Marc Anderson, Verisign.

TOBIAS SATTLER: Tobias Sattler, United Domains.

RICK WILHELM: Rick Wilhelm, Verisign.

[LEON MCFESSEN]: [Leon McFessen], 1&1 IONOS.

TOM LAM: Tom Lam from Cloudflare.

ROGER CARNEY: Roger Carney with GoDaddy. Thanks, Zoe, for taking care of us.

JODY KOLKER: Jody Kolker, GoDaddy.

RYAN [LONERGAN]: Ryan [Lonergan], Donuts registry.

JOHN RUPP: John Rupp, name.com.

[KATHERINE RUDINGER]: [Katherine Rudinger]. I'm Donuts Registry and name.com.

QUOC PHAM: Quoc Pham, Neustar.

ERIC ROKOBAUER: Eric Rokobauer, Endurance International Group family of registrars.

MARC ANDERSON: Thank you, everyone. For those of you in the back of the room, I won't make you come up, but if anybody wants to introduce themselves, please go ahead.

[DEET WALENDIN]: [Deet Walendin], Valideus and Com Laude Registrar.

MARC ANDERSON: Thank you, everyone. Appreciate that. Again, like I said, welcome to this meeting of the TechOps group. Zoe has an agenda for us up on the screen. We are in Zoom, if you want to follow along on there.

Looking at the agenda, we've got a session on Auth-Code management and security. Rick Wilhelm has bravely volunteered to take the first

topic for us this morning. And then, we'll go over to an update on transfers. Owen's agreed to give us an update on that. We'll talk a little bit about what's going on in the EPDP, both the Phase Two work and IRT, particularly as it relates to TechOps—the technical aspects of that that we may care about.

And lastly, we'll spend this morning's session on a little bit of future planning—looking ahead at how to get the most out of TechOps. The idea of this session is to try and get feedback for Tobias and I to consider planning agendas and topics for the next year of TechOps. So, at this session, we're sort of winding down the calendar year, so we'll be looking ahead to the next year—what kind of things we want to talk about and do in TechOps—how to make this a valuable use of everybody's time.

I think when we started this off, it grew out of the GDD Summit, and some of the side conversations we had, and we decided there was really a need to keep momentum going between GDD Summits, and that's what this has turned into. At least for me, it's exceeded my expectations. I think this has been really useful. But I'd like to see this continue to be useful, and hopefully make this something that people want to spend time doing and working on. So, we'll wrap things up talking about that, but in the meantime, like I said, Rick will be kicking us off with a discussion on Auth-Code management and security.

RICK WILHELM: Very good. Thank you, Marc. We've got slides coming out of Zoe's email in a minute here, so we'll let those come up. While those are coming up,

and you're tormented by my face on the video screen up there, we'll just give a little bit of background on this.

This is related to an internet draft that Jim Gould and I, both from Verisign, currently have working in the IETF Registration Extensions Working Group. It's being pitched … I shouldn't say "pitched." It's being presented as a best practice—a BCP—as opposed to something which is a standard, which would be a requirement. This is something that's being presented as a best practice, that we would be recommending registries and registrars would use.

As you can see, being someone from a software background myself, I believe in reuse. You can see that these slides were repurposed from a presentation that Jim did at IETF 105, which is the previous meeting. So, we're just reusing these, since the essence of the paper hasn't really changed.

What we're going to talk about here … You can flip to the next slide, Zoe, please. Thank you. What we're going to talk about here is that the situation that we have is that the AuthInfo is something that everybody depends on—registries and registrars depend on—for being the mechanism used to allow transfers to fundamentally proceed.

This is not getting into how the registrar … There's a lot of discussion about the registrar, and the form of authority, and that sort of thing. This is separate from that. This is about how AuthInfos flow between registries and registrars, and how the registries and registrars generate and handle those things. It's that sort of a discussion.

So, it's an implementation detail that doesn't really relate to the IRTP—Inter-Registrar Transfer Process—that's talking about. It's subject to a PDP that's going to be coming up. I just want to frame it as being related to but distinct from the PDP work that's going on. This is complementary to, but doesn't collide with that work.

Now, of course, if for some reason that PDP would decide to eliminate the concept of the AuthInfo, and I have no idea whether it would, then this draft would obviously be, as it said, overcome by events. Does that make sense. I just want to frame the context as within the existing transfer policy, and the notion that AuthInfos are central to the way that transfers happen—not so much how the registrant goes about acquiring that AuthInfo, but how the AuthInfos are trafficked between the registry and the registrar, and the processes used for handling those—just to set the stage what we are and aren't talking about.

You can see that in the middle bullet here about what is out of scope, on the slide. There's a paper on this thing that no doubt someone will paste the link into the chat, but of course, the way that search engines work, it doesn't take much to pull this out of a search engine and find the link to the internet draft. It's relatively easy to understand for those that are within earshot of this discussion. It's not written in very dense prose or something like that.

By the way, I have five or six slides here. If you've got questions as we're going on, please feel free to raise your hand in the air. Someone can alert me if they see something in chat, and we can go into discussion on this. This isn't a carefully-prepared presentation or something like that.

So, let's talk about what we're doing here. The BCP is a collection of proposed best practices that registries and registrars should follow, and in no particular order here … You can flip the slide, Zoe, please. One of the things that we're advocating for is to make the AuthInfos be produced with strong randomization.

We give a mathematical definition of what is meant but "strong." We're recommending at least 128 bits of entropy. That means 20 characters-worth, if you're using all of the printable ASCII characters except space. There's a reason we don't include space, because one of the reasons is that the AuthInfos are transferred using EPP, and leading or trailing spaces are problematic in XML, as those of you that deal in that element of the protocol know. They get truncated, and so that would lead to havoc and mayhem. That would be bad, so we don't want that. So, we leave space out of the character set.

When you do the math, and you want to get to 128 bits of entropy, that puts you at 20 characters. That would be a recommended minimum. If a particular registry would want to set a policy for me characters, that would be within their right, but this is a recommended minimum. There's math that's in the paper that shows how that entropy calculation works.

So, that's one thing. This, of course, would keep approaches out, which would have things like the same AuthInfo for all the domains for maybe a whole registry, or for a registrar, from generating the same AuthInfo for all of the domains within its portfolio, or various variations on what

security people would consider “bad.” That’s one element. Make these things strong random, and it defines what “strong random” means.

The second thing here that we have is making these Authorization Infos short-lived. What do we mean by that? What that means is that the default state of AuthInfo in a domain that was not in the transfer process would be null. So, what that means is that under most of the domains in a registry, on a percentage basis, would have their AuthInfos at the resting position of null.

It’s only when the losing registrar would provision on AuthInfo into the domain, that then it would be non-null, and then that AuthInfo would exist in a non-null state, strong random, for a period of time that the quote unquote transfer window was held open, and then either the domain would presumably transfer, if the transfer was approved and completed by the gaining registrar, or the losing registrar would go through and say, “The TTL that I have internally set on this is done. The transfer window is close, and I’m nulling it out. If you need a new one, dear registrant, you’ve got to come back and get a new one.”

Now, registrars would have the ability to leave those provisioned in there for a relatively non-trivial period of time—two weeks, five days, seven days, up to registrar policy, whatever you want to do. My personal anticipation is that different registrars, if these kind of things were adopted, would have different policies, down to a per-registrant and a per-domain basis, perhaps even contextual on what’s going on in their business at that time.

For example, if Registrar A, coming from Customer B, on Domain C, at a particular time, and whatever's going on their account—and some registrars can have really sophisticated AI algorithms—it says, "Wait a second. There's something really weird." This domain goes over to special handling, and they get someone on the phone, and they say, "Your AuthInfo code is valid for an hour. Because this is a high-risk domain, it's a short whatever." Other ones, they says, "You know what, look, this domain and this customer … We're not really that worried about it. Your transfer code is good for a week—" something like that. That might be the default.

But what it does is the default state is that it's null, and then the AuthInfo code gets provisioned for a short period of time. Kind of make sense? It's sort of like when you get a login thing on your phone, from when you're trying to log into your bank. It's not two-factor authentication, but it's sort of like that.

Next item—storing the AuthInfo securely. The registrar client should not really be storing that authorization info, and the registry server side should store it using a cryptographic hash. In other words, don't store it in plain text. That would be a bad thing to do. The client side—the losing registrar—after they provision it, after they create it, should only hold onto it long enough to give it to the registrant. They shouldn't squirrel it away, because if they've got it, then it could be leaked, or stolen, or whatever.

And the server, when they store it, they store it as a hash, so that when get it from the gaining registrar, they can do a subsequent hash, match

it, compare, say “thumbs up” or “thumbs down,” but they shouldn’t store it in plain text. It seems obvious, but again, we’re just documenting some things that are good practice.

Perhaps the thing that might be a little bit interesting or surprising is that the registrar is not storing the thing, which means you really can’t give it back to the customer, which is kind of an interesting thing. So, if you didn’t give it back to the customer immediately, what you would do would be you would generate a new one, re-provision it, and then give it back to the customer. But the thing doesn’t sit at rest on the client side. Okay, next slide, please. Yes, Quoc?

QUOC PHAM: If you’re talking about storage as well, if there’s logs involved—because obviously, you’re sending things through a VPP—then it’s probably another suggestion to make sure that you don’t store the logs, especially the AuthInfo, in plain text as well. Remove that information altogether, log-wise. Just a little bit more than storage.

RICK WILHELM: Ben, please.

BEN MCILWAIN: Seeing as how we’re going to be mandating the amount of entropy on the tokens themselves, should we also mandate some aspects of the way that we’re storing the hashed token? If somebody just sees the unsalted MD5, that wouldn’t be very good. Do we want to say it has to

be at least so many rounds of so-and-so good algorithm, and it's salted as well, maybe?

RICK WILHELM: Thanks, Ben. That's good. The logging on the storage, we might cover that. I'll have to reread to remember if we do that. The recommendation on the storage of the hash, I'm not sure if we have a recommendation on that. I don't think we do. I think that's a good one to include. Let me dig into that. Thank you. That's a good one.

ANTHONY EDEN: Just looking the current version of the spec, it does say that the plain text version for the authorization information must not be written to any logs by the registrar or the registry. So, that piece is already covered.

RICK WILHELM: Very good. And there might be a way to strengthen that a little bit, to your point, because I could imagine, if I can crawl inside of your head, you might say, "Yeah," and you might be able to say, "Look, don't log it, even in any form, unless it's really necessary," or something like that. But we'll look at that regardless, so thank you. And thanks for the check on the spec real quick.

That's a description of what the items are. Let's look at a … This is just on the screen. I've sort of talked through this, but this puts it into text

so you can see it. Let's look at what a transfer flow using this sort of an approach would look like.

At the beginning of the registration cycle is where number one sits in. This is not the transfer cycle, but at the beginning of the registration cycle, the registrant goes to register the domain with the registrar, and the registrar sends in the create command in number two—this is all at the beginning—with an empty or null AuthInfo over to the registry. And then, the registration gets created, presumably successfully. We'll take the sunny day path here. We won't get into rainy day stuff.

And then, so it's sitting … The registration is sitting the registry database for some period of time, and then subsequently, later, presumably after all of the other conditions are met sometime in the future, the registrant wants to transfer the name. Setting aside exactly how the transfer policy would get the registrant to this point, they now request the authorization info from the losing registrar.

So then, what the losing registrar does in number four … There's a big time gap between number two and number three, just to be clear. In number four, the losing registrar, at that point, generates the AuthInfo value in a secure and random fashion, sends it to registry over EPP, and then, without storing it, the registrant would give it to the registrar.

And then, also in four, although it doesn't say it here on the slide … If I were better at doing these slides, it would, but then, maybe it wouldn't fit on the slide in the right font. The registry has to store it securely. So, the registry would store it using the hash—using the description that

Ben provided. And then, the gaining registrar would come back and verify the authorization info using the info command.

This is assuming that this is what the transfer policy says. That's why it says "optionally" there, because this is a mechanism for policy, and not to be memorialized in a BCP like this. This is the approval step on the transfer.

If the transfer policy would require it, number five is where the verification would happen—where the gaining registrar would do the verification. And then, the gaining registrar would … What the gaining registrar is doing there is making sure that one exists, and then the gaining registrar sends the authorization info on the transfer request command. It gets thumbs up/thumbs down verified by the registry, and then, upon the successful transfer, the registry would null the old AuthInfo out, and then the domain returns to its at-rest state post transfer, of having an empty AuthInfo.

That brings us … Basically, the domain would end up in its lifecycle, somewhere between step two and step three, as those of you in this room and within earshot presumably understand. It's how that gets put into practice. Here. Any questions about that. Yes, Ben?

BEN MCILWAIN: I just have a small question about step five. So, any registrar will have the ability to request the information on any domain name, regardless of who it's owned by, and see whether or not there is a transfer token set, and that's all they need for that?

RICK WILHELM: This BCP doesn't change that mechanism. It doesn't change the access provided to the AuthInfo. So, whatever the registry policy is around that availability, it doesn't change that mechanism. Whatever rights that the registrar has to get that information is not changed. I'm pretty sure that there's not a way to determine if the AuthInfo has been set. Let me see. I think that Jim Gould is also attending our meeting here. Jim, are you in chat? Could you hear Ben's question? Let me see if Jim's in chat.

BEN MCILWAIN: Absent the ability to see if one is checked, I'm not really sure how number five would work, then. How would it actually check it, because they're not the registrar of record, so they certainly can't see the authorization [inaudible].

RICK WILHELM: This BCP isn't changing the access to the AuthInfo data. Jim says, "We explicitly did not allow to determine if …" Sorry, I'm incorrect. "We explicitly did not allow it to determine if the AuthInfo is set." Yeah, so this doesn't allow … Jim, are you on audio at all. Let me see if Jim's here. He might not be on audio. No audio? Okay. Jim, if you want to type some stuff into chat around step five, that would be helpful.

ANTHONY EDEN: For anyone who's interested, in 4.3 of the spec, it talks about this. I was thinking one thing that might help here is show the two response types

that would come back, based on success or failure, because right now, it essentially shows the client request sending the password, but it doesn't show what the EPP response is going to be. That might help clarify this. It does explicitly say that it must not return indication, though, of whether the authorization info is set or unset.

RICK WILHELM: Very good. Thank you. Jim says that patching the matching Auth-Code will indicate that it's good. That's in six. That's good feedback, Jim. This doesn't change it. Please go ahead.

VLAD DINCULESCU: Just two points. On point number five, the "optionally…" How I understand it is that if you provide an AuthInfo as part of a domain info request, and that AuthInfo is incorrect, you don't get back the domain information pertaining to the domain name. There's your verification process right there. You get a 2202 code—whatever it is—authentication, authorization error. Essentially that, I think, is why we're doing number five.

And then, what I want to know is with regards to number seven … You're looking at the tie-in to where the losing registrar updates the AuthInfo, with the tie-in to number seven, where the registry sunsets it. Those are two different policy perspectives over there. Have you considered a point where a registry would have its own internal policy—a global policy for its operators—whereby, after x amount of time, if the registrar hasn't updated the AuthInfo based on its internal policies,

then it automatically gets unset by the registry, be it five days, ten days, however long, at a registry level?

RICK WILHELM: Good question. We did consider that, and we didn't go that way, primarily because it gets into very … This ends up, in these certain situations, getting the registry into the middle of the transfer process, because it will inevitably end up where there will be timing situations, where it turns into registry policy, and the registry will be interfering with the transfer, because the registry is wiping out an Auth-Code that a registrant is about to use, that the registrar had intended to be lasting a particular amount of time.

Also, it would encourage … I shouldn't say "would encourage." It could lead to divergent implementations among registries about what their timings are, and also might encourage registries to have varying policies per domain type. And so, it introduces a whole layer of complexity, and also would, for the first time we think, get registries even deeper into having a policy that could break—I shouldn't say "break—" interfere with, or get in between transfers that are going on between two registrars.

And so, the easiest and simplest thing to do is to not have the registry unset existing AuthInfo codes. Now, a registry may choose to do something to issue a report to registrars to say, "Did you know you have AuthInfo codes that are over 15 days old—" something like that— "over 30 days old—" something like that. And registrar may choose to do that.

Bad practice, exposes their users, but it's really not necessarily any different than a registrar not setting transfer lock on the domains.

So, that's why that thing isn't in there. That is one where we made a very deliberate choice, after considering some options. We looked at the second- and third-order implications of what that looks like. At first glance, it does sort of look attractive, that, "Oh, well let's have a timer at the reg …" But then, you get into what happens about different registrars, and what registrars might want to do, and it leads to a rat's nest of implications. Does that help? Yeah, please.

TOM KELLER: Will there be an option for registrars to unset the Auth-Code, because that is completely missing that picture. You set it, and then it's there forever.

RICK WILHELM: No, sorry. The registrars do have the option to unset the Auth-Code. The registrar that is the current sponsoring registrar does have the option to unset the Auth-Code at any time. You can change the Auth-Code and either provision it or set it to null. I would expect that registrars will have different policies related to different domains, different customers—all sorts of different situations that they will allow Auth-Codes to have widely and wildly-varying Auth-Code times, as short as 10 minutes, and longer than 10 days is where I would … There's going to be an interesting bell curve of that. Quoc?

QUOC PHAM: So, the registry unsets the authorization information. I assume that means the ability to nullify AuthInfo for a domain, making it untransferable. Especially for just the general domain create process, you have to provide that on create. So, there's an implied—for me anyway—meaning that there has to be an AuthInfo with a domain at all times. So, this unsetting of an AuthInfo, effectively nullifying a mandatory field, changes it a little bit from the policy?

RICK WILHELM: No. The AuthInfo code can be created empty, and the protocol does support that, to be able to create it empty, with an empty AuthInfo code on create.

QUOC PHAM: Is the RFC … I don't have the schema in front of me, but is it not a token that has to be at least one to 32 characters or something like that?

RICK WILHELM: It can be not provided. It's okay.

QUOC PHAM: On create?

RICK WILHELM: Yeah. You can create it with no AuthInfo. And in this case, what the registry does here is it only unsets it upon a successful transfer. So, what

that does is that it wipes out the value that was provided by the losing registrar, and would set it to null then, only upon successful transfer.

QUOC PHAM: Just to add to that, I think, just off the top of my head, with regards to the schema, I believe that domain AuthInfo, domain PW is a requirement. The text value in PW can be null. It can be an empty string entirely.

ANTHONY EDEN: I'm looking at 5731, and the AuthInfo element definitely is required. I'm trying to look through here. The challenge here is also … I'm always wary of nullifying, versus empty stringing, versus leaving out completely. There's a lot of room for misinterpretation there, so it's a risk.

RICK WILHELM: Yeah. Zoe, could you flip to the next slide. Then Marc can keep bothering you. Here, we've got a slide that does talk about this, and the RFC support. I think right here it's a little bit of a lead-in. It allows this mechanism … It's supported to be able to provide both the updating … The first bullet item is on create. The second item is to allow the currently-sponsoring registrar to wipe it out with an update command. So, if it's previously been provisioned, it can then be wiped out. So, that works within the protocol.

QUOC PHAM: I guess where I was getting wrapped up was with one of the first slides, saying that an AuthInfo has to be at minimum 20 characters. I guess you can say that if you do provide an AuthInfo, it has to be a minimum of 20 characters.

RICK WILHELM: Well put. Yeah, good clarification. Thank you. And then, at the bottom, we talk a little bit about the info response. This gets to some of the … Basically, what we've done with the way that this has been set up with the BCP … The basic gist of this slide here, slide five, is that the BCP is set up to be implementable within the existing RFCs, and not require any changes to those existing RFCs. That was one of the main goals of writing the BCP—that it could be implementable without going in and changing anything to there.

You can flip to the last slide, Zoe. The last slide just gives an overview and a conclusion here—so, no EPP extension. Only set it during the transfer process. It can have a client-managed TTL, but it doesn't have to. And with the recommendations, it's not stored by the registrar, and stored as a hash by the registry.

The mailing list that's discussed here is the IETF RegExt mailing list. I'll be going over these slides, although as not as much detail, during the RegExt portion of our meeting. We'll probably rip through these in about five minutes during that thing. I wanted to have a little bit deeper discussion here in TechOps. The folks at RegExt have, like I said, heard a bunch of this before.

But this is really the principal audience for this, because this is a BCP that, while it was written in the IETF, it's really directed at this group here, and something that we think would be a set of best practices, that if implemented broadly, would generally help improve the security of the handling of AuthInfo within the Registry/Registrar community, which could be something of a general improvement, in order to increase the security of the transfer process, and the security of domain name registrations in general for all concerned.

So, I'm not sure what our time is looking like, Marc, or if anybody has any other questions. I've gotten some good notes here for some good feedback, and Jim Gould is also on the line, and been listening and taking notes also. We'll be working on incorporating some of these into our next version of the document.

At Verisign, we're going to be working on—have this work underway. But obviously we're looking for, to the extent that registrars and other registries are interested in doing this … This is not a standard. This is like some other things that you're doing, that you'd be implementing, but it's more of a set of best practices, like I said. We'll be discussing this in IETF. Any other questions, comments, or thoughts? Yes, please.

TOM KELLER: How does it tie into all the other things we've been doing around transfer? This is pretty much along the lines of the white paper the TechOps group put out. Can we link that to it? Is there any further discussion planned? We're now having the scoping team, and they will

probably have that end up actually pointing to a document produced by TechOps. This sounds beautifully going into the right direction.

So, is there already some larger Registry support behind that? Have you guys talked about it? This is the first time I hear about it as a registrar, but how can we get that into the process. It's not we have end up with two competing documents that say the same thing. That would be a bit silly. So, how do you want to move forward from here. That would be my interest.

RICK WILHELM: Anybody else want to comment? I don't need to be the first to respond. Going one, going twice.

ANTHONY EDEN: Would it be possible for us to endorse these best practices, rather than having to write it, if we agree that from a technical implementation standpoint they meet our expectations? It seems like it's just one piece of the rest of the work around transfer. Maybe an endorsement of the BCP might be a good way to do it.

VOLKER GREIMANN: I like this new process, but it seems to be at least skirting some of the policy requirements currently in place. For example, the FOA, even though the Temp Spec been extended, I think ultimately, we will have to have a policy replacement there. Is there any plan to start the policy development process on this? How to ensure that this process comes

through that unscathed might also be a good question that we should ask ourselves before we launch that.

RICK WILHELM: Very good. Tom, did you want to …

TOM KELLER: Sure. I think that's pretty easy. There's a group now set up to scope the transfer policy—a PDP. So, from my understanding, this group has to convene in two weeks, and then we will probably go into the exercise of coming up … What parts of the transfer policy as existing, and what we want to solve first. It think this is where we can point to this prospectus and say, "Okay, this has already been debated by the tech folks, and this is how we should implement it in a certain way."

That's nothing to do with all the other implications about FOA. This is a policy-related question. It has to be solved in the PDP. I think the PDP will be forthcoming, after the scoping team decided how the scope should actually look like.

RICK WILHELM: Jody?

JODY KOLKER: Volker, I was just curious. You said that this skirted the FOA, but I don't see that, so I'm just curious. It's just about passwords. It's not about

getting the authorization, gaining FOA, or anything like that. I'm just looking for clarification.

VOLKER GREIMANN: The way that the process looks to me, it doesn't require any FOAs anymore. Unless we left that out in the presentation, and still require those, the FOA process is still very much in the policy, and still very much enforced by ICANN Compliance. It would either have to be supplemented by that, or we would have to have the policy development process taking out the FOA.

I think if that process is starting now, that's going to be very good, but I think we have to be very cautious of other interests try to amend this into a more complicated monstrum that we might not like or recognize from what we've seen now.

JODY KOLKER: Yeah, from what I saw on this, it's basically just the interaction between the registries and the registrars. It doesn't really deal with any of the registrants at all, because we're just dealing with the password updates. That's the way I saw this, and I was trying … To me, it looks fairly contained, so that we don't have to get into all of the other stuff yet. We'll save that for the policy development process, which will be fun, I'm sure.

RICK WILHELM: Correct. All good comments. This very explicitly did not address the FOA. Jim and I tried very hard to keep it compartmentalized around the AuthInfo codes. Our thinking was as follows. We started this, actually … We started sketching ideas in a notebook, actually, before the PDP work was announced. The first draft of this was a while ago. I won't say a long time ago.

Then, when the PDP got announced, one of the things we realized was that this actually might be very helpful to the Contracted Party House, because what it would do could help to compartmentalize the discussion within the PDP to say, "Look, when it comes to the AuthInfo codes, if the process is going to use AuthInfo codes, we've got security around that—some best practices. We, the CPH, knows how to handle AuthInfo codes."

And so, if the process is going to include AuthInfo codes, it doesn't need to be a debate about how they're going to be securely handled, transmitted, supported, stored, and all that stuff. Therefore, it could help encourage the positive development of a good discussion around the FOA, or not FOA, or how that all turns out, by helping to keep that discussion from exploding into a big, proverbial dust cloud that starts to include the FOA, and the Auth-Codes, and this, and that, and, "Ooh, we ought to …"

The transcript is going to be a disaster right there, so I feel really bad for … But those of you that are watching the video or here in the room, you get the idea that Jim and I hoped that it would help to say, "Look, from a technical perspective, if it's going to include AuthInfo codes, we, the

technical community, know how to handle those." We wanted the technical community to contribute some thoughts, and take some good feedback here from others in the room, which we've gotten, which we really appreciate, and we'll incorporate that, and then allow the policy development process to tackle the stickier policy problems—but sort of to compartmentalize it. Is that helpful?

So, yeah. It explicitly doesn't take on the FOA, but it tries to let the FOA discussion focus on the FOA, and say, "Look, if you're going to use AuthInfo codes, we got that—" we, being the CPH team.

MARC ANDERSON: Thank you, Rick, and thanks, everybody, for the discussion there. Auth-Codes transfer process is something we've talked about in previous meetings, so great to continue that conversation and get everybody's thoughts and feedback on that.

Next, we're going to turn it over to Owen. Owen, hopefully we didn't step on your territory too much here. But Owen's agreed to give us an update on the status of transfers themselves—so, a little bit of any easy transition there from one to the other. So, Owen, over to you.

OWEN SMIGELSKI: Thanks, Marc. No problem. I understand the abbreviated nature. Even though this is a passion of mine, I can certainly cede some time, although I do have to point out first, I'm really impressed by the stickers on your laptop. You've got AC/DC and Avril Lavigne, so you're really doing all ends of the spectrum there on music. I really appreciate that.

I'll be real quick here about transfers. There's two big issues going on. First is the gaining registrar FOA, which TechOps identified before GDPR became effective in the Temp Spec, that there'd be some concerns with that. Tried to engage ICANN Org in some discussions about what we could do about that—how that would be impossible post-GDPR Temp Spec. That fell on deaf ears, and then of course we had the problem of not being able to do the gaining registrar FOA due to a number of issues—lack of consent to process the information, email addresses being unavailable because they weren't functioning, etc.

So, this has been ongoing with ICANN Compliance for a little while. Namecheap in particular has one where basically Compliance has told us, "You need to implement this," which kind of came to a head in Marrakesh, where the registrars … We met with GDD and some other representatives from ICANN Org.

What has resulted from that is on October 9th, the Registrars Stakeholder Group sent a letter to the GNSO Council. I'll paste the link into chat, if anybody wants to see that. That includes a brief explanation of the problem, and requests that the GNSO Council write to the ICANN Board to get a deferral on the gaining registrar FOA, pending the outcome of a PDP for that.

It's been discussed in Council at the October 24th meeting, and I know it was sent to the GNSO Council email list, although there's been little activity since. So, I don't really know what to expect out of that. Hopefully we'll have some more progress on that. There was also … It's

not included in that, but we also sent a draft letter for the GNSO Council to send to the Board, so hopefully that can progress relatively quickly.

Also, Tom touched on this earlier about the Transfer Policy Review Scoping Team, so I won't go into too much detail on that. Basically, we're going to be coming up with what we think needs to be discussed to fix in the transfer area. There's a link there to the working group, and there you can see it's a couple of registries, a lot of registrars, and somebody from NCUC, who have expressed interest in that. We are trying to meet in the next week or so, after ICANN 66. That hasn't been set definitively.

I wish I could give a timeline on when we think it will be completed. We had hoped and internally discussed and of year, maybe early next year, so that's why there's a question mark on the target completion for that. But that will then feed back to the GNSO Council to formally launch a PDP for that. That's pretty much it. I don't know if anyone has any questions. I can certainly answer them.

GREG DIBIASE: Just so we're aligned, the best-case scenario here is that Compliance defers and this is settled in the transfer PDP, and that's something that once we have the scoping team, we can go back to Compliance say, "Look, we're already figuring this out," basically?

OWEN SMIGELSKI: Correct. Yes. What will actually … The GNSO Council will send a little to the Board. It will be similar to what happened with the enabling and

disabling privacy proxy under the change of registrant, where the Board sent a letter, basically directing Compliance to defer enforcement. So, that would be, we'd imagine, something similar along that lines, and through the open ticket that Namecheap has, I'll be updating Compliance. We'll also do that here ICANN 66, to let them know about what we're doing—that we're working our best to get rid of that requirement or defer it.

MARC ANDERSON: Thank you, Owen. I appreciate you stepping in and doing that. What Owen doesn't know is this was an interview for him, and if he did a good job on that update, we'll ask him to keep doing updates on that. Thank you, Owen. Appreciate that.

OWEN SMIGELSKI: Appreciate being voluntold.

MARC ANDERSON: Before we move on, I know a bunch of people in TechOps signed up for the scoping team. Just out of curiosity, how many people here are on the … Excellent. Pretty good representation there. I look forward to continued updates on that one, and hearing how that goes. I think our next agenda item … Speaking of voluntold, Sarah has been kind enough to jump in, to give us an update on where the EPDP is, both the IRT and the Phase Two EPDP work. So, over to you, Sarah.

SARAH WYLD: Super. Thank you. Good morning, everyone. Yes, Marc invited me to do this about 10 minutes ago. Thank you. So, the EPDP has completed Phase One and issued a set of recommendations. Those are now in the implementation phase. The way that process works is that ICANN has an Implementation Team. They take each recommendation and figure out what actually needs to happen to make it effective in real life. And that then gets presented to the members of the IRT to review as a section of what will eventually be a full policy document.

So, we've taken each recommendation, turned it into a piece of policy—a little chunk of policy—and discussed as a team, "Do we agree that this is the right approach? Should we change the approach?" Sometimes it's very straightforward, and we all agree quite simply, "Yes, this is how we should do it." Sometimes, it is more controversial, or there's a lot of back and forth discussion. But we've been a pretty effective group at coming to agreement on how we should handle those things.

We're almost completed working through all of the different recommendations, but we do not yet have a full document of the new policy. We were hoping to have that here in Montreal, but it looks like that's not quite ready yet. So, instead, I think what we'll be doing at our meetings on Wednesday and Thursday morning is to review the final recommendations, the policies coming out of them, and get everything ready, so that we can then come to look at one single policy that we can see as a whole and in context.

The original goal for our policy to be released was actually back in August, so that we could have our six-month implementation buffer

with effectiveness at the end of February in this year. Unfortunately, we did not hit that date. We all worked very hard, and work as fast as we can, but there is just so much to be done. Also, not every recommendation is quite final.

So, for example, my favorite, Recommendation 12, is about the organization field and how that's handled. The Board didn't adopt the whole recommendation. They had some concerns with some of it. So now, we in the Implementation Team cannot finalize how to implement that until those concerns are resolved. Even when we're done going through all the recommendations, we don't quite have the full policy document ready until all the recs are done, so that's where we are. So, the date for which this will be required is not yet set, but it will not be February 29th. That was the IRT in five minutes.

Then, on to Phase Two of the EPDP. That is focusing on developing a standardized system for access and disclosure of nonpublic registration data. Sometimes I wake up in the middle of the night, and I find myself saying that phrase—"standardized system for access and disclosure." This is a remarkably complicated process, with many different people having a lot of strong opinions as to why data should be disclosed or should be not disclosed, under what circumstances, for what purposes, and to whom.

It's really complicated. It's a really exciting time, and a great example of this multistakeholder model, and how we can get so many different people to all come together and share what they think is the best thing to do for the future of this system. They've made some really good

progress over the course of the day yesterday. It was a grueling, 10-hour day meeting yesterday, so thank you to any EPDP members who are here, because that was intense.

I think they came up with a final accreditation building block. It's divided up into the building blocks—each of the different pieces we need for this overall system. There's accreditation. There's acceptable use. There's which pieces of information should be included in which circumstances. The team is just working through those right now. I think that remains on track. Nothing is ever as quick as we want it to be, but also these are really complex topics, and you can't rush things. So, that's where that is. That was a very quick recap, but I'm happy to take any questions. Or not take questions, but defer them to Marc.

ROGER CARNEY: There was a letter sent recently on a huge stumbling block from the group. Is there any update on that letter, as to if ICANN is planning to respond sometime soon?

SARAH WYLD: Right. So, this is the letter to the Board saying, "We really need to understand what role ICANN will take in this process." I'm simplifying it. In my own head, I know what I think the answer should be, but I certainly haven't heard anything back about, or I don't know where that is.

MARC ANDERSON: Thanks. Great question. We did get an update from—I think it was Dan Halloran—yesterday that they're working on it. Hopefully we'll hear soon, but we don't have a timeframe on when we'll get an answer.

RICK WILHELM: Real quick question, Sarah. I think I know the answer to this, but I'll ask it anyway. February 29th, 2020 won't be the date, but is there a so-called date for a date?

SARAH WYLD: No. It would be nice if there were. One thing we don't have, really, is a set project plan with timing on it. That's not how this particular IRT has been set up. So, when we have the whole policy, then we can … Eventually, it will be set, and then we can set our six-month buffer period, so then we'll know the date, but no, we don't have that.

ROGER CARNEY: The six months was the minimum, and the IRT believes that it will probably be longer than six months needed for implementation, especially if we have to have consent ready before this goes out. That may take a lot more time. So, the six months is a minimum window of notification from ICANN to Contracted Parties. The Implementation Team needs to look at what this final policy says, and determine how much work needs to be done, which will mostly likely be longer than six months. Thanks.

SARAH WYLD: Good point, Roger. Thank you.

MARC ANDERSON: Thank you. And Sarah, I appreciate you letting me put you on spot there. Just another couple quick shows of hands again. Show of hands if you're on the IRT, in the room. And then, if you're on the EPDP, member or alternate. So, there are a number of us in TechOps that are participating in both the IRT and EPDP.

If you have questions, feedback, input, please feel free to snag any of us between sessions, or any time you can track us down. Obviously, we're very interested in the technical aspects, in making sure we produce something that's implementable at the end of this process. So, any feedback in that area is always appreciated.

I think this brings us to our last agenda item of the morning session. I mentioned this at the top. We want to take this opportunity to solicit feedback from everybody in the room on what you want to see out of TechOps in the coming year. I'm just going to throw this out there, and hopefully people will be willing to come to the mic and talk about what you've liked, what you haven't liked, what you'd like to see, maybe what you don't want to see.

I'm just going to leave this open-ended and solicit your input on this one. I'll just give a time check though. We're coming up on a 10:15 break. So, we have a little bit of time, and I'll just start it off, and see if anybody wants to come to the mic and be the first brave soul to jump in the queue.

ROGER CARNEY: I think this group has been very good at collaborating. Even if we don't agree on things, we can agree not to agree on them. But I think the group works really well. I think we've talked about it in the past, that the challenge for this group is what do we do with what we did? We do a lot of things, and produce some things, but we don't know how to get that to the next level.

I think that we see, technically … We already just saw an example of Verisign taking this to the IETF and moving something that's been talked about, and technically moving that forward. So, I think we have a technical path in moving things forward. I think the policy side is the stumbling block that we run into. If we have a good idea, how do we move that?

Owen showing this letter actually helped me take maybe the next step here. How do we take what we think are good policies inside of this group and move them forward? Owen showed us where the Registrars Stakeholder Group requested the GNSO to do something. Maybe that is a path this group can take, is we can post these. If we come up with something we agree to, we can take them to our stakeholder groups, and then they can push that forward through that.

One of the things I … Again, we've talked about this in past meetings. I don't want to get too much administrative and bureaucracy in this group, because I think we'll lose a lot of what we get done. So, I'd like to use some infrastructure that's already built if we can. Just my thoughts. Thanks.

MARC ANDERSON: Thanks. Tom, I know you've shepherded the transfer discussion for a long time. I don't know if you have anything you'd like to add to that.

TOM KELLER: Yeah, definitely. What we see and experience is that there was a lot of good work done around transfers, and we came to a stumbling block, not knowing what to do. One of the findings we had to discuss in that that whole TechOps exercise is not officially existing in any way, shape, or form. The question is whether we cannot or we should not formalize it.

The problem we have in ICANN, that there is only policy, full stop. What we do is more like technical standards, and trying to figure out ways of how we can work together as registries and registrars, and it's not binding. The question is whether we really need any kind of policy process to agree on certain things, or we can set up an own regulation around best practices we agree on—some process, whatever.

I'm totally with Roger, saying we shouldn't overengineer the whole exercise. But with what we see currently, it's really hard to come from results to implementation, and I think before we engage into further discussion about things we want to do, we should have that discussion first—how we want to move forward in the future to make things happening. If we don't figure that one out, we can come up with all brilliant ideas, but it will not change our world.

MARC ANDERSON: Thanks, Tom. I'll throw out a couple other leading questions. We've had a couple different formats at ICANN meetings. We've done sessions where we've had guest presenters over the course of a day. Also, we've done breakout sessions. We've done whiteboarding, brainstorming sessions. And we've done a mix of the two. I'd love feedback on what you'd like to see at our face-to-face meetings—which of those worked, which of those don't work for you. And also, what you'd like to see during our biweekly calls—how we can make that a good use of everyone's time.

RICK WILHELM: Thanks, Marc. On the transfer policy, I have a … I don't know if it's a different view. I view the transfer situation as a bit of a success, in that I think that without the work that the TechOps group had done, I don't think that it would have been possible to get a PDP started in that amount of time, with this level and credibility, and with this kind of focus.

I would really encourage folks in CPH TechOps to view the initiation of a PDP around the IRTP as a clear win for TechOps, in that the PDPs in this kind of area just don't get started every weekend. And for one to be focused on work that this group, and a topic that this group started, is actually pretty important and pretty profound, and it's something that the group hasn't accomplished before.

When we were at GDD Bangkok, and we were sitting in that very warm, very little room, and Pam was like, "You all have to do a PDP on this," and a bunch of people in the room went, "Ugh," I was like, "Good.

There's an outlet for it." And it was a positive outcome. Now the PDP has started. We saw Owen's slide up. There's a scoping team—this sort of thing.

I think it's really a positive that this group, setting aside all the caveats … I recognize it's been a while in coming, but it's really a positive thing that this group was responsible for initiating that on such an important topic, that's been so central to what has been going on at ICANN for years, and years, and years. So, I would view it as a clear win for the group. Thank you.

GREG DIBIASE: Yeah. I was thinking the same thing, but adding another element of that. It's a clear win that we're starting a group with a pretty clear roadmap, or at last an idea of where we want to get to. We'll see how this works out, but I'm cautiously optimistic that now that we have not only what our plan is, but justifications behind it—like, why is the Auth-Code strong enough? Well, we've done all this work. I view that as a really big one.

JIM GALVIN: Thanks, Marc. I want to make an important distinction, I think. I agree with what Rick was saying. I think that it's absolutely essential. We regard it as a success. I think there was a bit of a small success in front of getting to this PDP, and that as the Temp Spec that was written, because we had to make some really quick changes and agree to that in order to get even the Temporary Specification put together so we had

something that would work. And we did that, so we have had some significant success along the way.

What I want to build on is there's an important distinction to be made here. I think it's fair to say a PDP process is a fairly heavy-weight process. Domain transfers probably falls into the category of something which really warrants that heavy-weight process, in part because the existing one is a consensus policy, and those are created an managed through PDPs, so you're kind of stuck with that.

I really think the distinction that Thomas is making is is there a category of more lightweight things, that it might be useful to find a process, or a path by which we could come to some agreement about them. We have this ad hoc, if you will, as compared to a PDP process, best practices domains website, where there's a bunch of stuff listed there. It's all interesting, and I think that a lot of that stuff is going to move forward. We're going to have some discussions about that this afternoon in the TechOps going meeting with RegExt.

But it is a fair question to ask of ourselves. What is a process? What is the path by which we might do more lightweight things that are not consensus policy? Such a thing does not currently really exist in ICANN, effectively. Maybe there's a role for TechOps to play there at the technical level, and we just have to brainstorm and think about what the right way is to make that happen. Thanks.

MARC ANDERSON: Thanks. Volker?

VOLKER GREIMANN: I think we have a very good situation here. We have basically what I would call a win already, because this is one of the very few PDPs where we are pitching. We're not on the receiving end. We're not playing defense. We have initiated something that we see necessary, that we want, that we're driving the process at this moment. This is something that is very rare within ICANN. I would see that as a win already.

To go beyond that, to your question, I think we should ask ourselves the question, "What do we want to be?" Do we want to be a formal element within ICANN that is similar to the SSAC or ALAC, that is an advisory body to ICANN, that has powers to initiate PDPs or provide advice to the Board? Or do we want to be what we are right now—a more loose federation of contracted parties that pitch ideas and try to get them into the GNSO and the policy making process, and influence policy that way.

I think both sides need to be looked at, and I think we have to make up our minds what we want to be. Being more formal is, I think, I a lot of work, because it requires some changes to the ICANN structure, but it might be worthwhile. I'm just not sure what we want to be.

MARC ANDERSON: Thanks, Volker. I'm going to go to Graeme here.

GRAEME BUNTON: Good morning, everybody. Graeme Bunton from Tucows. Also SG Chair for Registrars. Great discussion. Jim, I think you made really good points about … Yeah, transfers is a really big deal. It belongs in a PDP. I would think it unfortunate if all of the good output from this group needs to go through that process. I don't think that's the right answer.

A solution might be that the GNSO figures out how to do more than one type of PDP, so that there is some sort of tech-initiated, smaller, more lightweight, technical-focused initiative that isn't everyone and their dog in there, and nothing gets done, and it takes three years, and it's really frustrating. How much room there is in the current PDP 3.0 process for that right now, I don't know, because I haven't been paying a lot of attention to that. Maybe I should.

But it could also be that we just figure out ways to do things amongst ourselves that don't require anything like a PDP, because it's maybe just easier if we figure out the lightest-weight thing. There should be some category of problems that fit into this, where we can just go, "Guys, we all agree that this is a better way to work to implement whatever feature. Let's just go do that." Now, there's no carrot and there's no stick, other than our lives get better and easier, and the technology improves. I would love to talk more about that in some other forum, to figure out how we can move forward there. Thanks.

MARC ANDERSON: Thanks. I know I have a couple people wanting to jump in. I'll just note, we're at our break. So, if you guys could be real quick, I'll just remind you, you're standing between everybody else and coffee.

VLAD DINCULESCU: That just sounds like a horrible thing.

RICK WILHELM: I'll be quite quick. While I'm empathetic to the point that Jim is making, I would caution this group with trying to touch the policy process, because it can have a lot of unintended second- and third-order effects, because it wouldn't necessarily just apply to what stuff that TechOps is doing, sort of as Graeme was implying with the discussion around PDP 3.0. So, that's probably about enough.

VLAD DINCULESCU: Just a quick point about what Graeme said … A lot of the stuff that comes out of here is essentially best practices, and things that we can all cumulatively agree this is a good approach to doing something. It doesn't have to be instilled within a PDP, and the majority of these things essentially can be implemented without the need to circumvent current policy.

Look at transfers. Look at this Auth-Code management, who's to say that this Auth-Code management goes against any form of policy? It doesn't. It doesn't touch on FOA. It doesn't touch on the requirement for a sponsoring registrar to provide the AuthInfo code as easy as possible, as to provide anything else. This is just simply an add-on that makes life easier for a lot of people. A lot of the work that we do here is simply just that. It's a lot of good-faith work that comes out with a

whole bunch of easier implementations and much easier work for everybody.

MARC ANDERSON: Thank you. I'm going to have to draw a line on this conversation, but thank you, everybody, for the feedback and the discussion. Quick check … We have a 15-minute break, and then we'll be picking back up with a meeting of the RDAP Working Group. So, if you enjoyed hearing Rick's voice, come back in 15 minutes. You'll get to hear him speaking some more.

RICK WILHELM: And if you don't …

MARC ANDERSON: You've been warned. So, thank you very much, everyone, and hopefully we'll see more of you throughout the session.

ZOE BONYTHON: Thanks. We can pause the recording, or stop, yes. Thanks.

**[END OF TRANSCRIPTION]**