ICANN73 | Virtual Community Forum – GNSO: CPH TechOps Meeting
Thursday, March 10, 2022 – 10:30 to 12:00 AST

DEVAN REED: Hello and welcome to the Contracted Part House TechOps session. Please note this session is being recorded and is governed by the ICANN standards of behavior.

During this session, questions or comments submitted in chat will only be read aloud if put in the proper form as noted in the chat. If you would like to ask your questions or make your comment verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor. Please state your name for the record and speak clearly at a reasonable pace. Mute your microphone when you are done speaking.

This session includes automated real-time transcription. Please note this transcript is not official or authoritative. To view the real-time transcription, click on the closed-caption button in the Zoom toolbar.

To ensure transparency of participation in ICANN's multi-stakeholder model, we ask that you sign in to Zoom sessions using your full name—for example, a first name and last name or surname. You may be removed from the session if you do not sign in using your full name.

With that, I had the floor over to Marc Anderson.

MARC ANDERSON: Hello, everyone. This is Marc Anderson. Welcome to the CPH TechOps meeting at ICANN73. On behalf of my Co-Chair, Jothan Frakes, I would like to welcome everybody to this meeting. You can tell it's the last day of ICANN meetings because Devan there seems to be losing her voice a little bit. So hopefully you're doing okay to run the home stretch.

By way of a little bit of introduction, I want to talk a little bit about what TechOps and welcome everybody to this call. TechOps is a joint venture between the Registry and Registrar Stakeholder Groups within the Contracted Party House. We might generally once a month as our current cadence. And our meetings are typically closed to Registry and Registrar Stakeholder Group members only. However, for ICANN73, we made the decision to have a public open session of the TechOps group.

And so, with that in mind, first I want to welcome anybody who is joining TechOps for the first time. Hopefully, the discussion about Oreos and cookies to start off the meeting gave you a hint of what kind of fun you're in for today on this TechOps call.

To give you a little bit more background on what this group is and what we do and how we formed, you really have to go back to the early GDD Summits, the first couple GDD Summits, for the Contracted Party House. During those summits, there's a lot of excitement and enthusiasm, but we were sort of unsure initially what to make of those GDD Summits. We had open sessions where people could suggest topics and have breakout sessions and discussions based on the topics that were voted on that had the most interest from attendees.

And from those early GDD Summits, we realized there was a lot of interest in talking about technical topics. The GDD Summits had a fairly large participation from members of the Registry and Registrar Stakeholder Groups that were less focused on policy and more focused on technology and the technical issues and challenges facing registries and registrars. And so we actually had dedicated technical tracks on subsequent GDD Summits.

But another trend came up. And what we realized is that the GDD Summits were occurring once a year. And those discussions we had during those yearly GDD Summits were great, and we made some good progress, but we were losing momentum in between. We realized that, when we met for subsequent GDD Summits and were suggesting topics to consider, we were picking up the same topics from the previous year and no progress had been made between the GDD Summits.

And so, from that, the idea was put forward. And I give a lot of credit to Graeme Bunton, the registrar chair at the time, for championing the idea of having a technical forum for registries and registrars to come together and talk about common issues and concerns and, in particular, talk about the technical touchpoints between registries and registrars, where they have common concerns and challenges to tackle.

So that was a little bit of the genesis.

And Zoe is pointing out that I forgot to mention Tobias. I believe he was the Vice-Chair of the Registrar Stakeholder Group and one of the founding vice-chairs of the TechOps group and was certainly

instrumental in the creation of the group and contributing to its early success.

So hopefully this was a fun trip down memory lane for those of you who've been members and participating in TechOps. But for those of you who might not be familiar with TechOps, what we are, and what we do, hopefully this provided a nice primer to the group and what we're all about.

Let me just pause there. Any questions? Any comments anybody wants to make before we move on.

Seeing no hands, why don't we jump right into it? And maybe, actually, before we do that, go back a slide. I'll just tee up the agenda real quick. Okay, thank you. So by way of agenda, at previous TechOps meetings, we discussed ideas for topics to cover at this meeting. As a reminder, typically our meetings are an hour. This meeting is 90 minutes, so we have a little more time to work with for today's meeting. And we're going to cover some future work output. Jothan Frakes is going to walk us through some of the ideas for future work we're considering, as well as looking at ongoing PDPs, working groups, and how that may affect the TechOps group or bring up topics that we may consider.

Rick Wilhelm is the Chair of the RDAP Working Group and agreed to give us an update on what's going on in the RDAP Working Group. Jim Galvin, as his tradition, will be providing an IETF, specifically on the REGEXT Working Group. There's a lot of overlap between topics that happen there and topics that get discussed here in TechOps. And Jim,

a Co-Chair of the REGEXT, has really been doing a great job bringing the two groups together and fostering communication and coordination.

Roger Carney will give us an overview from the Transfer Review Policy PDP. He's the Chair of that working group. By way of a little background for everybody, the transfer policy has been something of particular interest to the TechOps group. It's something we've talked a lot about. And so this is a topic near and dear our TechOps, if you will. So we're looking forward to hearing from Roger on that one.

And, lastly, Sarah and I are going to talk a little bit about the SSAD ODP and the subsequent Operational Design Assessment (ODA) that came out of that. This is a topic that we've been wanting to take on within TechOps for a while. And with the recent publication of the ODA, the timing seemed good to bring it up. In particular, if the SSAD were to proceed and be developed, contracted parties are going to have to figure out how to integrate with that SSAD system. And so this is a topic that we've been following very closely and have been wanting to have a more substantive discussion on. And like I said, the timing seemed right to capture that here ICANN73.

So that's what we're looking at as a high-level agenda.

And with that, I will turn it over to Jothan Frakes, who will lead us through our first topic. Over to you, Jothan.

JOTHAN FRAKES: Thank you very much, Marc. Great introduction. And, definitely, it is a real privilege to co-chair this with you, I just wanted to say.

So maybe just to plus-up the background and talk a bit more—if you'd advance the slide—I'm going to be talking a little bit about our future output. But in order to talk about our future work, which is driven by recommendations and suggestions, it's really important to acknowledge some of the work we have been doing and talk about that in the context of why it's so important that the registry and registrar technical operational engineering teams are having these conversations. It's really to tie into understanding how to make our industry work better from a technical sense, how to create standardizations, how to address technical implementation challenges in areas where … For those of you who are familiar with the OSI's seven layers of technology, there's a joke that there's layers 8, 9, and 10 above that, which are the users and maybe money, policy, governance and other things. And we address the seven layers.

And a lot of what happens inside of ICANN happens outside of those very technically-focused things. And what we do is we end up having to implement a lot of that technically in simple terms. And so the effort that we put together is to follow those different policy development processes outside of what we're doing and really try to find out areas and give early feedback as to what aspects of those may or may not be technically reasonably accomplished, or perhaps there's an affection of policy that might cause something unexpected to happen somewhere

else due to all the interdependent parts, al the gears of this machine of EPP, SRS, and our delightful registration industry.

So when this group founded, it was largely an output product of the 2012 [round]. There was quite a lot of new top-level domains that were introduced through the new TLD program. And it introduced quite a lot of different new entrants as far as technology goes with implementations of SRS systems. And there were things introduced, such as premium name tiers. There's was DNSSEC requirements. There was all kinds of new things that had to happen in our industry.

And many different providers would implement them in many different ways. And this, for registrars, was quite challenging because, where you had already attenuated to maybe two or three or maybe four providers and were able to technically integrate into them, you now had to do that for many more. So it expanded the z-axis of complexity in how many places you need to write integrations. And each of those different technical providers that were registry backends may have implemented them in different ways.

And so what we, I think, as registrars, really speaking from the registrar perspective, benefit from in the work of CPH TechOps is that we're all sitting down together in a room and able to talk through how others are implementing things in order to make sure that, as we propose or do things, it's easier from a registrar standpoint and more attractive to integrate, but it also follows a standard behavior that you'll be able to provide a consistent experience to your customers, the registrants.

So some of the things that we've tackled when we get together … We meet monthly currently, but we often would meet face-to-face and gather and sit down and have a whiteboard. And people would provide as to things that we really … What are some pain points? What are some challenges? What are some areas that we need to discuss and work towards standardizations on? And some of those things have come in the form of following the outputs from different policy development processes—their parts—of ICANN and understanding that we as contracted parties will have obligated things that we need to implement.

So currently what we are doing is following the PDP work, things that are happening in the implementation view teams, and also following things happening outside of ICANN in the IETF standardization, which Jim Galvin will be talking about shortly.

But we have had some topics that come through that have been proposed, and we really don't sit and proclaim what those topics are. They will often evolve from needs that come in the business world or things that come in helping to improve the stability, security, and resilience of the system or just make it less complex for different people to implement.

So some of the things are standardization and structure of registry maintenance notices. Registrars would be identifying, "We'd really like some consistency so that, when a registry needs to do their scheduled maintenance or unscheduled maintenance, there would be some

uniformity to how that noticed would be received so that we can plan and somehow action that."

There are other things such as EPP registry premium fees. Those didn't exist prior to the 2012 round. And so as those were implemented, there was a variety of different ways that the registries had implemented them. And we worked together to create standards that would help to unify so that it was a more consistent experience across registrars and thus for registrants.

And that's not a comprehensive list. Those are just some basic examples of those types of topics that we come through and work with.

And what happens is that someone will propose this and then we'll get together a group of people to socialize that idea, to talk it through. And often that will form into proposing RFCs to the IETF to create standards. And we do that collaboratively, talking either as registrars or registries or even both nowadays to talk through those different items.

Right now, we're having conversations in and around a lot of standardization of RDAP. And you'll be hearing from some of my other colleagues about other areas that that's developing.

So it's not a comprehensive list but, as we talk about what we're going to doing as future work output, we're constantly pulling and evolving this.

Our current efforts are very heavily focused on things happening inside of the policy development work at this moment with SSAD, as we'll be hearing later, the Registration Data IRT, the work of the RDAP Working

Group, things related to data accuracy, and the Transfer Policy Review Working Group. We will be seeing things that need updating or changes. We may see modifications to contracts.

And so this group, the Contracted Party House TechOps, in summary, is the place that we have those discussions. We, inside of the Registrar Stakeholder Group, have a poll open currently for registrars to submit topics. And we're already receiving topic submissions that we'll bring to the table. And there is really no closed period for when to submit these to the group. But we do often, in the presence of an abundance of those, vote them up or down to help understand the places that we can most focus. Again, we're at an unprecedented time with a lot of different things happening inside of the policy development process, so our efforts have really been attenuated to being ready for those and providing the input into this different areas. And you'll be hearing from each of those different groups following my discussion here about our future work output.

So for registrars in attendance, we have a form that's made available to the Registrar Stakeholder Group where you can submit topics or concepts. And we have an ongoing list of different items. I mentioned just a few here today, but we often will talk about the rotation of SSL certificates and how we communicate with registries. We may talk about certain IDN extensions. It's often very, very focused on the registration or resolution processes and how we can work on constantly improving that. So we look forward to submissions. We're very open and inclusive about taking on those. And that's a little bit of an insight into our future work and what we'll be doing.

But we currently spend a lot of time helping to identify areas that require some technical specifications or adjustments inside of other groups. And you'll hear a little bit about that as my colleagues from within the group as they give their various individual group updates will be telling you about.

So are there any questions about our future work output?

Hearing none, I'll go ahead and pass it on to our next speaker, who is Rick Wilhelm with PIR, who's going to talk a bit about RDAP. Thank you.

RICK WILHELM: Very good. Let's talk a little bit about RDAP. RDAP, of course, is the—you can flip to the next slide, Zoe, please—protocol that's replacing WHOIS for publishing registration data at registrars and registries. The RDAP Working Group is a group of technical experts that work on the RDAP profile. And that's the document that supplements the specifications that are in the IETF RFTs that describe the requirements for registries and registrars that are under ICANN contract and how they need to implement RDAP as an output mechanism. It may be interesting to ccTLDs but it's not a contractual requirement, unless they have some other contractual obligation implemented. But this is particularized for ICANN, the contractors.

So right now, the work that the RDAP Working Group is working on is updating the RDAP profile documents. There are two of them. They come in a pair—a document called the Technical Implementation Guide and the Response Profile. And that work is being done to

incorporate the output of the EPDP IRT (Expedited Policy Development Process Implementation Review Team, for those of you that are not facile in all the ICANN acronyms). Those documents were published previously as the first version coincident with the publication of the temporary specification. And they needed to get updated for the output of the EPDP.

Most of the changes are related, as it says here, to the differences in the output of some contact objects. So, for example, the admin contact is expected to be going away. The technical contact is getting a change of shape. It's getting a little bit less data—some things like that. So the RDAP output has to adopt to that.

One of the key contributions that's also being incorporated into this version is a mechanism that is going to be incorporated into the RDAP profile that is used to signal when data is being redacted. The temporary specification requires the output to include the string redacted for privacy, spelled out just like you see it there on the slide. That has issues if the output format requires something like an e-mail address because, as everybody within earshot knows, "redactedforprivacy," written like it is there on the slide, is not a valid e-mail address. It wouldn't pass. It obviously has spaces and it doesn't have an @ sign.

So what we've done is coordinated with the IETF REGEXT Group. And you'll hear more about this in Jim Galvin's section. There's a draft that's been worked up and published. There's a link to it there. Or you can just use your favorite search engine on "REGEXT RDAP redacted" and be

able to find it. And this uses an RDAP-compatible mechanism to signal that the field is being redacted using the RDAP output format rather than using the actual text in the field.

Other changes that are happening is that the RDAP profile is going to incorporate some implementation flexibility regarding an emerging technical implementation called JSContact, which stand for JavaScript contact, which is a replacement for jCard. Technical implementors don't really find jCard difficult to work with. And so this document right here, which you can see a link to, makes that easier using JSContact. So the proposed changes to the profile will allow implementers to have the option to implement either JSContact or jCard.

And so the working group is undertaking this work now, and the profile documents are expected to be ready for public comment as part of the implementation review team's package with all the policy output that goes out for public comment. So the group meets biweekly, generally, unless something comes up. And the work is generally progressing well.

Happy to take any questions if there are any. Thanks.

MARC ANDERSON: Thanks, Rick. Not seeing any hands. Before we move on to the next section, I do want to say there's been a lot of activity recently within the RDAP Working Group. There's a flurry of activity as we get those profile documents updated. As Rick said, they're changing from the temporary specification to the upcoming registration data policy as the driving policy for the profile. And so making all those changes has led to a lot of

activity within the group. So there's much more excitement than we've had in a while in that group.

Also, I'd be remiss if I didn't thank Rick for the excellent job he does chairing that group and keeping them on track and moving forward, to which you get a quick +1 and exclamation mark from Sarah. So not just +1 but a +1 with an exclamation mark!

So with that, if we can move on to our next topic. And here I'll turn the microphone over to Jim Galvin from Donuts, who will give us our IETF REGEXT update. Over to you, Jim.

JIM GALVIN: Thank you, Marc, and thank you, Rick. You very carefully and lovely covered a couple of the things that will be on the list of things that I'll talk about here. But it's all good. This is the benefit of working together so well. I speak here today as the Co-Chair of the IETF REGEXT Group. Antoin Verschuren is my Co-Chair, but he's not with us here today.

Next slide, please. By way of introduction, I think it's very helpful for this group to understand really the relationship and what REGEXT is. REGEXT stands for REGistration EXTensions protocols working group. And the IETF is an organization which has over 100 working groups, and each working group has its own particular scope.

The REGEXT Working Group has a scope that basically includes three things. As we all know, there are two protocols that mean a great deal to registries and registrars. A lot of what we do—pretty much everything that we do—is dependent on these two things, and those are EPP and

RDAP. Rick just told you a little bit RDAP and some of the things that are going on in that particular working group, but those two protocols represent two things that fall into the scope of REGEXT. So anything that's EPP- or RDAP-related would fall into our work in the IETF. And the third topic area is anything which is related to a technical specification of a relationship between registries and registrars or any other technical activity that they might make.

One example I'll give you from history is, of course, the escrow specifications. gTLDs and registries and registrars of course are obligated to escrow data on behalf of the system as a whole. And the technical specifications for how one accomplishes that and what those look like are actually IETF specifications. And they would fall within scope of the REGEXT Working Group as it is today.

The IETF may highlight a couple of things here about joining this REGEXT group. It's open to anybody and everybody. In the IETF, they try to stress that you participate as an individual. And the only thing that's really required is mailing list participation. You don't actually have to come to meetings. So you don't have to be physically anywhere. In fact, it is a rule within the IETF processes that all decisions are actually made on the mailing list. So even though there are meetings, and people come and they contribute, that's the way in which work gets done. So I do want to encourage you. As we say in every TechOps meeting, we're always looking for people to join REGEXT. I do want to encourage you to come along and join.

I also want to mention the data tracker. If you've never heard of the data tracker in the IETF, I encourage you, really, if you're interested in technical standards, to get yourself a DataTracker account. The most useful feature of the data tracker—it has a lot of really good features; trust me—is that you can go to any document and you can diff that document against any prior version in its history. There is a tool built into the data tracker that's there that let's you take anything. And this is an easy way for you to see the differences. So even if you haven't been paying attention for a little while, and you see a document that's about ready to come and it's going to be done and you want to know how it's different from a prior version, there's a diff tool for that. And it'll show you all those differences and highlight them for you. So it's really quite valuable.

Next slide, please. And on this slide, this is a list of the currently active documents within the REGEXT Working Group. And I'm going to step down through these, just a little bit, and highlight them for you and hopefully pique your interest and make you want to join the REGEXT group and come and be a part of discussions and recognize them. There is actually an effort to create an EPP extension for exchanging internationalized e-mail addresses. Obviously in ICANN, we have quite a problem in universal acceptance. And one of those elements is obviously being able to use an internationalized e-mail address. And it turns out that we need a way to make sure that those can be exchanged in EPP in all of their glory.

And so here's the technical specification for how to do that. And so if this is something that's important to you, it would be good to you to

check this one and look at it. It's actually just about ready to be submitted for publication. So it will leave the working group and be submitted to what's called the IESG, the steering group, where then it gets brought to an IETF review and it goes through kind of like a public comment period in ICANN. For those of you who are used to in that terminology, think of it in those terms. And then it'll proceed on to the publication step.

The JSContact is one that Rick already talked about. As you know, in RDAP we do take advantage of the jCard specification. And this is a proposal for how to migrate to using JSContact, which, from an implementation point of view, implementors tend to regard JSContact as a little bit easier to deal with. It works a little nicer with internationalization. Your mileage may vary.

An important part of this specification is being able to migrate from one to the other. So the whole transition process is actually documented in there, too. But it's there so that now we have that option and something that we can do so folks can think about what they want to do there and make that available to their RDAP clients.

As Rick also said, redacted has become an interesting thing. In the old WHOIS days, since responses were strictly straight up a text dump of just a bunch at you, one of the most important features of RDAP in general is the fact that it's structured data. That's what you get back. The response is not a pretty print text dump. It's actually structured data so that clients can then do whatever it is they need to do in order

to better serve their users, including translation and transliteration, for example.

But one of the interesting things is that, as redaction became something important in policies within ICANN, then it became essential to really that, oh, RDAP did not have a way to signal all of this, didn't have a way to manage all of this. So, as Rick said, the RDAP Working Group is dealing with the technical profile and implementation guide for RDAP implementers on the ICANN side. And what the REGEXT Working Group has been doing is looking at how to provide that signal inside the baseline RDAP protocol so that it's useful for all of the various systems that use RDAP, not just domain names.

The next one down is the search capabilities. One of the interesting things from the 2012 round, of course, is that there were some searchabilities. And many registries and registrars has implemented those and added them to their WHOIS implementations. That was actually a part of the review process back then.

And so the interesting thing is that, in RDAP, the baseline specification has a very limited, small set of options for search—certainly, a pretty far cry from the various searching capabilities that were available with others.

And I have to ask a question here. I don't seem to have the screen in front of me anymore. I can keep talking about is it just me or is there something going on technically?

JOTHAN FRAKES: Zoe, did you open up a window in front of … Sure.

JIM GALVIN: Okay, now I see you. Now I see me. Maybe she'll just restart the display of the slides.

DEVAN REED: Hi, everyone. It looks like Zoe might have dropped. Let me just try to get in touch with her a different way.

JIM GALVIN: Okay. I'm going to keep talking. Il speak to the slide. I do have an alternate copy in front of me. It'll be fine.

So, search capabilities. And so this document is a specification for how to a broader set of searching options at RDAP. It's useful to keep in mind that RDAP is used not just by domain names. It's also used by the number community on their side of things. So the specification does need broader review than just in ICANN. And so this is just a look at how to search on other kinds of elements in the data that's available. This document, too, is pretty much ready to be submitted for publication. So if you think you are interested in search, now is kind of your last opportunity to dig in and have a look at this document and make sure it's something that meets with your needs and actually makes sense for something that you want to implement.

So the remaining three documents are newer bits … Well, no. So the next two were more recent work. We have a new document which has

just been adopted: the registration data dictionary document. This is an interesting document. It's being used to create a central location—a registry, if you will—of all potential data elements that might fit into a registration system.

It's interesting to observe that, in fact, there is no central authority or central source for registration data elements. Some would argue that the contracts that gTLDs have and the specification that's in them in terms of what data they collect is passed on or the data that they display in the WHOIS specifications that are an essential part of that are kind of an authority. But that's just not centralized. And there are some variations among registries. And of course, the numbers community has its own look at these kinds of things. So there are a lot of different places where these things might exist.

This registration data dictionary is an attempt to create an IANA registry that can be used on a first-come, first-serve basis to define new elements and define existing elements. And so, for right now, they're drawing in a collection, trying to reach out and pull in various technical specifications for data elements so that we can have a central location where these things exist and have them all defined.

Note, by the way, that this is … I want to draw a really hard line between his registry existing and defining elements and their technical specifications from what you do or don't have to provide when you respond in an RDAP response or in fact what a registry or registrar might collect and have available for other reasons. Those are all policy issues and they're covered elsewhere. They're not part of this specification.

This is just a collection of all of them and putting that out there. So that's a good thing.

The next document down is the simple registration reporting document. And going back to some of the history that Marc and Jothan were offering us, this document actually has its origins this TechOps group. We had quite a discussion a couple of years ago in one of our TechOps meetings about the relationship between registrars and registries. In particular, registries and registrars—no , Slide 11, the next one, please—exchange data. Registries ordinarily provide reporting, all kinds of documents, to registrars—for example, a list of transactions for the month and various other kinds of reporting things that they need in order to continue their business and of course to reconcile the relationship between them.

So there was a point in time when we had a big discussion and there were actually nine different documents that were proposed, one for each different kind of report that could be used between registries and registrars. This is simply a specification that takes those nine different reports and creates a reporting registry, again, at IANA that allows you to indicate standard reports and any other general report that you want. But it does indicate a mechanism for creating those reporting files, making this concrete. It's basically an indication to use a CSV format for creating the documents that you send and the reports that you use. It's just a way in which to do that if people want to adopt it.

Part of the discussion in talking about all of these things back in the day was just that registries and registrars … There was no standardized

mechanism for doing that. So registrars often found themselves having to process all kinds of formats, ranging from nice mechanical ones that were good for computers, as in CSV files or PDF files or sometimes other strange formats. So this is a way to create all of that into a standardized format, simplifying the way in which registries and registrars communicate.

And the last item is the federated authentication for RDAP using OpenID. This document has been around for a while. It came into existence when the RDAP Working Group first came into existence. It's a technical specification for how to do authenticated RDAP queries. Honestly, it's just sitting on the shelf right now in REGEXT. What it's waiting for is motivation which will come from policy discussions that get us to a place where we need authenticated queries. The assumption is that the SSAD and everything that's going on there will drive the need for authenticated queries. And this document will rise to the top of the pile in needing attention because it's expected to the protocol underneath supporting a portion of the SSAD and what it's going to do. But that's a discussion that we'll have when we get there. So for now, it's just hanging out there, waiting for an opportunity so that we can start to build against it and see if it's going to solve our problem.

So those are all of the current documents that are out there and available.

Next slide, please. I'm going to end just on this slide—so a bookend here, if you will. We had Rick talking about the RDAP Working Group, and I just wanted to end here with all of these links handy and ready for

you. There are actually five specifications that define the RDAP standard. In the IETF terms, it's called Standard 95 (STD 95), and if you go that link that's there for STD 95, you will actually see the five documents that define RDAP. And RDAP is being promoted. It's in the process of being promoted to a full Internet standard. All documents in the IETF start out as a proposed standard.

Of the five documents, four of them have actually already been promoted. They've had some changes to them—two of them; the 7480, 7481 (and it used to be 7482, 7483, and 74820), which were the RDAP specifications. The first two were just promoted. And this is where the data tracker actually comes in really useful because 82 and 83 did have some minor changes to the document before they were promoted to Internet standard—not material, mind you. You're not allowed to make material changes to a technical specification when you're upgrading from a proposed standard to a full standard. But there are some minor things in there.

And the last one of the list is the 7484. It's currently in the publication process. It's all the way at the RFC editor, and it's likely to come out real soon now. We like to say in the IETF, "Publication is imminent," and oftentimes that means it's six to nine months out. But this is at the end of that six-to-nine-month window. So it really is about to be there.

But using the data tracker, you can actually go back and look these documents up in the data tracker, and then you can compare them to the old versions. So if you want to see exactly what might have changed

from the old to the new, it's a very quick and handy way to go do that so you can check your implementation.

And I think that's it from me. I might have gotten us back some time here, I think. Happy to take any questions or comments. Please do come and join the IETF list and REGEXT. We're always looking for additional technical advice and insight. It's not just about any particular community—gTLDs or the names community. ccTLDs participates there, too. So if you're using these protocols or have an interest in them, we really could use some additional insight. And that's it from me.

JOTHAN FRAKES: Thank you, Jim. I may have spoken over Marc there. Thank you very much for covering that.

MARC ANDERSON: Yeah. Echoing what Jothan said, thank you for covering that. For my money, one of the nicest perks of TechOps is getting Jim Galvin's updates on IETF. Well worth the price of membership.

I also want to give a shot-out to Jim on the work he does in coordinating between the TechOps group and the REGEXT group. Jim has been a proponent of a couple of joint sessions between the REGEXT Working Group and the TechOps Group in the past. In fact, we've done that during some previous face-to-face ICANN meetings. And those have been well-attended and well-received. And so thank you for that, Jim. I appreciate that. And I think that's one of the big benefits of this group.

Looking at the agenda, we had planned to get our update from the Transfer Working Group. Looking at the list, it does not look like Roger has been able to join us, unless I'm missing him on the list. So if we could, Devan, I think that means we will skip ahead to the SSAD ODP slides. I know Roger had an overlap in commitment but was planning to join us. He was planning on running late, but we'll try and buy him a little bit more time by jumping ahead in the agenda to the SSAD ODP/ODA presentation.

And here, Sarah and I are going to take turns on this one. I asked Sarah Wyld from Tucows to give us a little bit of an overview on the SSAD ODP and ODA and hopefully explain what those acronyms are, what they mean, and give us a little of a primer for this topic.

So if I could, I'd like to hand it over to you now, Sarah.

SARAH WYLD: Thank you, Marc. Hi, everybody. Okay, get ready for some alphabet soup. All right, so as many of us likely know, the SSAD is a proposed Standardized System for access and disclosure—so that's our SSAD—of gTLD registration data. So parties which want to request that non-public data to be disclosed to them first would get accredited. And then they submit a request which is distributed to the appropriate registrar who reviews and makes a decision and either discloses the data or doesn't. so this was recommended in the second phase of the EPDP—this system. And these recommendations are now with the Board for evaluation.

But there's also this very special new Operational Design Phase. That's our ODP. And during this phase, the ICANN Org produced an assessment. So that's the Operational Design Assessment, which is a big report. And that report looks like it might cost to build and operate this SSAD in order to help inform the Board's decision.

So now the GNSO Council has also brought together a small team made up of councilors and former EPDP team members with the purpose of reviewing that ODA to give feedback to the council on whether it properly understood the recommendations, if there were incorrect assumptions or mistakes, anything missing—that type of thing. So that work with that small team is in progress. The small team will work with the council and then will meet informally with Board members to talk through that feedback. And these meetings are taking place over the next two weeks, with a final response back to the council by the 30th of March. So that's where we are with the small team looking at the ODA.

And there is some really interesting ideas happening in the small team for what the Board could do other than a flat-out rejection or approval of all the recommendations. So some members of this small team have suggested a pilot. And my concern there is that a pilot suggests to me that the whole system gets built and then only a small group of people use it. I'm not really sure what that would teach us. I'm all for agile development, building a small MVP, and then iterating, but I'm not certain that a mini-SSAD would answer our questions.

So the biggest unanswered question right now, even with the ODA, is still, what would the volume of use be? And a mini-SSAD with only a few

users won't answer that. So there's another suggestion, which I must admit comes from the registrar team that I'm part of it. And this suggestion is to use the existing ICANN ticketing system as the central intake point for disclosure requests. There's already a structure for taking in queries and distributing them to the appropriate contracted party and for tracking the response time. So this would let us track the full picture of request volume over the period of time, filling in the missing information about potential SSAD volume.

So there is our overview of what is the SSAD, what is the OPD, what is the ODA, and where it is right now. From our technical perspective here, I think there's really two big relevant pieces to think about right now. So number one is, how will contracted parties connect and interact with the SSAD? And then number two is, how is the permitted data actually distributed to the requester?

So I'm going to leave number one for Marc to get to, but for number two, I do just want to mention that there was a really interesting difference between the recommendations and what we saw in that ODA, the assessment. So in the recommendations, the expectation was that the requester will use what's called the central gateway as a central system to both submit their request and, if approved, to receive that registration data. But in the ODA, the assumption is that requesters submit it through the accreditation authority, which would pass it on to the central gateway. And then if they're approved, the requester then goes back to the registrar and says, "Hey, I was approved. Give me the data," and the registrar sends it them via RDAP, which is very different. And interesting difference.

All right, so that's my whirlwind review of what is the SSAD, what is the ODA, what is the ODP, and where is it right now, and from questions. I will pass it back to Marc, and I look forward to discussion. Thank you.

MARC ANDERSON: Thank you, Sarah. I appreciate you teeing that up. And I think, for full disclosure, Sarah and I are both recovering former members of the working group that recommended the SSAD and are current members of that GSNO Council small team. Yes, I'm optimistic that I'm recovering. I'm reading Sarah's chat.

As I said when I teed up this topic initially, what we've been wanting to do is have a little bit of a discussion about those touch points that Sarah just talked about. If the SSAD does indeed get built, contracted parties will have to figure out how to interact with the SSAD system and will need to understand and figure out how to deliver data (non-public gTLD registration data) to requesters for all approved requests, both of which will take a little bit of technical wrangling.

Now, normally, when you have a discussion like this, a specific outcome or goal would be in mind. I want to leave this one a little open-ended in that this discussion or deep dive could be just to provide more information to contracted parties on this call. It could be that we have a discussion here that leads to recommendations to a future implementation. And it could also be that this generates some follow-up and maybe flag some additional work or standards of specification that need to be done.

So I'm really leaving the purpose and outcome of this discussion a little bit open-ended. I don't have any preconceived or desired outcomes in mind. I just want to promote a discussion, raise awareness, and hopefully get people looking at this and aware of what it means, particularly from the Contracted Party House, considering what this would mean from a technical impacts and integration standpoint.

Devan, if you could jump to the next slide, please. So I'll start off by apologizing for the size of the text here. These are images from the ODA, the Operational Design Assignment, the report from the Operational Design Phase. And for our discussion, I did want to highlight both of these because they walk through the scenarios where contracted parties would need have touch points with the SSAD system.

And so the flow in the left is what would happen when a request is submitted for access to non-public registration data. And the contracted party would need to make a determination on what to do with that. And it goes through a number of decision points for how that flow would happen.

For example, there's steps for, what if the contracted party making the decision needs additional information? So there's a consideration for a feedback loop, if you will, where the contracted party responsible for that particular piece of data making the disclosure determination has to go back to the requester to get more information.

There's also, in this particular slide, a feedback loop for changing priority. The SSAD recommendations include the idea that there are different priorities of requests. And this includes a process if the

contracted party needs to change priorities of a request that's been received. Maybe something was submitted as a high priority but doesn't really qualify as a high-priority request or something was submitted as a regular priority but that priority has subsequently changed and needs to have some additional priority added to it.

So I thought this flowchart provides a good illustration of where those touchpoints would be.

What the system would ultimately look like, I think, is a bit of a question mark, but the contracted party would need a way to integrate with the SSAD system to perform all these functions. So the SSAD performs the routing of the request to the appropriate contracted party, but the contracted party will need a way to be able to see all the requests it has received, assess those requests, and take action. And I noted that there may be a need for additional data. So that feedback loop maybe needs to change priority.

And then ultimately though, the contracted party is expected to make a decision of either rejecting the disclosure requests, in which case there's an obligation to provide a reason for the rejection, or if it approves the request, [to] deliver the requested data to the requester.

And so all of that will have to be done through some kind of integration between the contracted party systems and the SSAD system. There's obviously different approaches to this. It could be done via a basic ticketing system. We all have accounts in a third-party application for tracking tickets. Or there could be some kind of robust API with

integration points for the contracted parties to help their own systems for integrating with the SSAD.

So let me pause here. This is probably a good time to stop and see if there are any thoughts, feedback, or discussion.

While I'm pausing, I'm trying to read chat. I'm not seeing any hands—

JOTHAN FRAKES: I've seen something in the chat from Werner Staub. "Has socialized an issue." Was that just a chat comment or was that a comment-comment for the session, Werner?

WERNER STAUB: Sorry. It's a comment on the subject here because I do we need an identifier and I think we do have capabilities as registries and registrars. And then there is the good intent of identifying actors. Actually, that's the business we're in. We've been doing this for 25 years or 30 years. So using domain names as identifiers for a specific purpose should be something that we think of, specifically because we have systems to manage the data, such as EPP. And then a domain name is a very good identifier because they can always point to [inaudible]. We can have a couple of constraints so as to avoid the [same party], which has come many times to apply [for] the different names each time. There's ways to deal with that.

We certainly should think of our own capabilities in managing data in the form of domain registries, actually, to do that. And essentially what

is the SSAD system if not a registry of parties who are supposed to be trusted to get access to other parties' data?

MARC ANDERSON: Thank you, Werner. Sarah, go ahead.

SARAH WYLD: Thank you. That is a really interesting and important thing to bring up. And I think you're very correct that the use of an existing ticketing system would not identify or verify that the requesters are who they say that they are, which was a big portion of the recommendations. There's a whole lot in there about the accreditation authority, governmental accreditation, and that would be missing from using a ticketing system in that way, yes, which is also one of the ways that it might reduce the costs of implementing the service.

For disclosure of data, it's very common that the requester of that data is affiliated with a government or with a business and is requesting data specifically in that capacity. So as well as verifying they are who they claim to be, the disclosing party would also want to verify their affiliation. Today, requests are sent directly to the relevant contracted party, and they would need to do whatever level of verification they deemed to be appropriate. So its not standardized but it does happen as any responsible data holder or data controller would need to make sure that you're disclosing to the correct person.

So there would need to be some type of structure for that, but I think, if we are doing a small-scale ticketing service rather than implementing

the entire SSAD, that would remain with the disclosing party rather than being a more formal standardized process. Thank you.

WERNER STRAUB: A question. How would you avoid coming out from hidden areas of the world? Maybe they're not actually who they say they are and, once they're identified as bad requesters, they just come again under a different identity. So you do need a way to identify them. So unless we create some infrastructure in the beginning, we probably cause, by o oversimplifying our problem, a big complexity of dealing with it.

So it might be a good use it. But specifically why I think it's domain names is that we have all kinds of rules already that are precisely the rules that the very people who want access to data invoke about what you should do in order to be accountable. So it is pretty good infrastructure already. We could have more of it on it but you could only add something if you have an identifier. Other than that, it just floats.

And identifiers must not proliferate, which is sadly what we have, which is a different problem. But the sad thing we have in the entire domain name industry is that contact handles proliferate uncontrollably. Nobody knows what they even have as contact handles. Every company has hundreds of thousands of contact handles and they don't know them. They have no idea what they [are] because they proliferate. So we have to find something that does not proliferate, that is under control. And that's a domain name.

JOTHAN FRAKES: And, Werner, thank you or bringing this idea in here. And I want to encourage you to join us for future CPH TechOps sessions where we go through and we litigate this level of detail and talk through things.

WERNER STRAUB: Okay.

JOTHAN FRAKES: I really appreciate that you brought that idea. And that's the kind of ideas that we have free flow for and that we discuss and hash out in the standard TechOps meetings. So I encourage you to bring it and join us at the next meeting.

Sarah, I don't know if you and Marc had been all the way through the topic. And pardon that intervention. And thank you for the comments. Yeah, "litigate" was probably a poor choice of words, but we hack things out and discuss them and socialize them. That's a better way to put it. Thank you.

And so, Marc, I hand it back to you.

MARC ANDERSON: Thanks, Jothan. +1. It's an interesting idea, Werner. I'd like to echo what Jothan said. I encourage you to join us in future meetings to hash that out further.

If we could go to the next slide, please. So before we dig into this one, I'll just note I see that Roger has been able to join us. So we'll cover this

slide briefly and then I'll hand it back over to Roger for an update from the Transfer Review Working Group.

So on this slide, this is an image from the Operational Design Assessment, the ODP report. And the reason I put this image in here is that I wanted to highlight the points Sarah as she was teeing up this discussion. And you can see here the requester flow and some of the differences here between what was envisioned by the working group in making SSAD recommendations. It's where the operational design assessment was produced.

In our working group's vision, I guess, we envisioned having the requester interact solely with the central gateway. And that's actually the reason for the name "central gateway." And rather than going through the accreditation authority, I'm not commenting on if it's a good idea, if it's a bad idea, if there's one better or not. I'm just noting differences here.

But also, as Sarah is pointing out, there's a two-step flow for the requester. The requester goes to the accreditation authority to submit their request and, once that request has been approved, they would have to take a separate action to a different location to submit an RDAP request to the contracted party. Here in particular—and I'm editorializing a little bit here—my concern is that one of the reasons for the SSAD, one of the advantages for the SSAD, is a single location to go to for all your request needs. And this flow would have the requester going not just to the accreditation authority but then having to go the individual contracted party to have their request fulfilled.

So I think this slide did a particularly good job of highlighting the flows in each of the operations or functions that would be performed between the different actors in the SSAD ecosystem and also very nicely bullets out all the functions or operations that we would have to perform as contracted parties, both integrating with the SSAD system and operating an RDAP server for the purposes of delivering requests. I note that, on the slide here, as they've documented it, there's also e-mail. So some information will be done via e-mail as well.

And I'm just noting in chat that Jan was also a working group member. And he's pointing that the central gateway … Originally accreditation and accreditation authorities were different roles, but we combined them. But also, I think an important point he's making is that, as we envisioned it, it was the central gateway that would have responsibility for connecting to the accreditation authority. And so it's the central gateway's job to interact with the accreditation authority in order to confirm the identity of requesters. So thanks for joining us and thanks for that point.

We're a little crunched for time. I'll give a quick call for questions or comments. Looking at chat, looking for hands. I'm not seeing any. I'll give a quick last call.

All right. With that, Roger, if I could turn it over to you for an update on the Transfer Working Group. And just a quick time check. We're looking at 13 minutes left in this session. So, Roger, can I hand it over to you?

ROGER CARNEY: Great. Thanks, Marc. Yeah, sure. It won't take me too long to get through this. As far as this group is concerned on transfers, I think the big thing probably is to know there definitely will be some changes coming and it'll be both—quite a bit of changes—for registrars and some for registries. So, again, when we get to the timeline stuff, I don't need people to be worried about immediately but definitely thinking about the changes. And, again, I don't think it's drastic changes we're looking at, but there definitely some system changes that'll need to be made.

Well, let's jump into this. Again, some of the big things that we've done is we're leading away from the term "auth-info" and going toward "TAC," which is the Transfer Authorization Code. And this is to be more specific. And I think, for auth-info, there was three or four different terms for that password—"auth code" and everything else that you can think of. So we wanted to be a little more specific and use TAC (Transfer Authorization Code) here.

And to be honest, I think there's at least a dozen specific recommendations around the TAC. Some of them are as simple as, okay, were' defining what the TAC is and what that means. But others are making sure that the code is going to be standardized across all registrars and registries. And, again, I think of "standardized" in air quotes a little bit because there'll be obviously some flexibility in the standards, but it will be much more standard than it is today, syntactically and things like that.

One of the other big changes is we're looking for the registries to actually write a one-way hash when the registrar stores it so that it is protected. So that's one big change for the registries.

And also, there'll be TTLs around that. And right now, there's a couple different thoughts on TTL, one being that there's a standard TTL. And I think the group came to 14 days. So the registry would need to actually enforce that so that, if a TAC was set, basically it's no longer valid after 14 days. But really the TTL is bigger than that, and the registrars can change that to a shorter period during that time. And "change that" is kind of in air quotes as well because that basically is just resetting it to be null or blank. To remove it would make that a shorter period if the registrar wanted it. Or if the registrant itself said, "Okay, I don't actually want to make this," on day 7, the registrar could go in and actually blank the TAC so, if someone tried to use it later, it'd be invalid.

Which brings me to the next thing. TACs will be single-use. So once a transfer is complete, that TAC and domain won't be valid any longer. So it is a one-time-use only.

And, again, there's more that we don't need to get into, but I just wanted for this group to recognize there will be changes coming, again, both on the registry and a lot on the registrar side.

Losing and gaining FOAs. We're getting rid of the to a degree. We're replacing them … I guess getting rid of the specifics of losing and gaining FOA had very targeted language that they had to be used and for a certain purpose. And we're getting rid of that and going back to just a prudent communication between the registrar and registrant. So

I think we ended up with several different spots of communication. And, again, I don't think that that has any effect on registrars, but registrars will need to start looking at that. And, again, looking at today's temporary spec, we've kind of gone away from these anyway. So we're just going to make it official and drop them and move to a better communication path.

We're finishing up some discussions on denial reasons—why a registrar can deny it or the opposite of that: why they can't actually stop a transfer. And we're going through that. And I think the big thing here for registrars is that there's some moving from that maybe you have the choice of denying but there's three of them that are actually moving to you must deny for at least three reasons. So I think it just provides a more consistent user aspect for the registrants.

And then I called this a transfer blackout period because there's been a lot of talk in several different groups about locking. And today there's optional lock periods, a domain create and domain transfer. And there are 60-day periods around those today. And the working group hasn't settled specifically. It's general agreement around a 10-day [inaudible]. And, again, I don't want to call it a lock because the lock doesn't really matter. It's just that they're not allowed to transfer a domain within 10-days of creation [inaudible] which for some technical reasons, the ATP invokes restrictions there.

In ICANN right now, we're trying to keep that the same for the transfer [and] post-transfer period as well, though we still need to have some discussions on a quick transfer reversal or what some people call

clawback. So, again, those are discussions that we're just finishing up. So we'll complete those. And, again, we're going to talk slightly about some of the clawback ideas, but we're not going to get into this in detail. We're just going to hit on enough to make sure that we're comfortable with the times we're setting.

Timelines. So the policy itself is broken into three different phases. Phase 1A is encompassing the true inter-registrar transfer policy. And 1B talks about the change of registrant transfer policy. So both of them are in one policy now, but it's two different sections of it. And Phase 2 is going to talk about all of the dispute mechanisms that are in place, including this quick reversal or clawback.

Timelines. We're looking at June of this year to release the initial report of Phase 1A. March of next year we're looking to do the initial report of 1B. And then those two will be combined into a final report to council because that'll encompass the whole entirety of the current transfer policy. And that's expected to go to council next August. And Phase 2 we actually haven't put any timelines on yet. We'll pick that up shortly after the final report or probably before that report goes and get going on transfer dispute mechanisms.

Hopefully that was quick enough and we got to the points that we needed. If there's any questions, please let me know.

MARC ANDERSON: Thank you, Roger. Not seeing any hands, I will see naturally attribute that to your excellent overview. Thank you for that. I appreciate you

joining and, as always, keeping us appraised of what's going on with the transfer policy.

I mentioned this a little bit at the start but there's some history with the transfer policy in this group. We've certainly been looking at and involved in that for quite a while now. And so it's excellent to see this work go forward. And hats off to Roger for the job he's doing chairing that group.

If we could just jump to the last slide, I think we're just about ready to wrap things up here. I guess I will, since the slide says so, I'll ask: if there is Any Other Business or any open questions anybody would want to ask? Now's your chance. Feel free. And I will keep talking while I look through hands.

First, I want to mention that there is ongoing discussion about the possibility of having a summit later this year. Instead of a GDD Summit, this will now be a GDS Summit. I think this is far from set in stone. It's, as I understand it, a possibility that's being discussed, and there's interest in doing that.

Assuming that does go forward, traditionally this has been a great opportunity for people in TechOps to meet face to face and have these discussions in person. In the past, we've often had a dedicated TechOps track at these summits. And a shot-out to Andee Hill, who has been very much involved in championing the summits in the past. So, Andee, I look forward to working with you again. Hopefully we can have a summit later this year. But the point there is, assuming we do have a summit, Jothan or I are committed to participating from a TechOps

perspective. And I'm hopeful that we can have a TechOps track at this, like we've done in the past. So watch this space as we move forward. Hopefully, we can do that in the future.

A quick housekeeping item. Our next TechOps meeting is scheduled for 13 April during our usual timeslot. We would have had a TechOps meeting next week, but this meeting will take the place of that. So if you still have a meeting on your calendar for next week, you can consider that cancelled. And I look forward to seeing everybody on 13 April.

Just in closing, I want to thank everybody that attended. We had around 45 participants over the course of this meeting. That's a great turnout. For those of you not familiar with TechOps, I hope this was a good introduction to the types of things we look at and discussions we have and what we do. So hopefully this was a worthwhile session for everybody. As always, the success and value of TechOps depends on the people attending and participating. So thank you to everybody that did so.

Jothan, any last words?

JOTHAN FRAKES: Well, I just want to thank all our attendees for joining us to hear a bit about what we do in CPH TechOps. And we are glad to have members of the Registry Stakeholder Group or the Registrar Stakeholder Group. If you want to participate and join us, we would be delighted to have you there. We welcome you. Just contact your Secretariat for participation information in joining. And thank you all for participating

today. It's been a real privilege to have you're here and to be able to co-chair this with Marc and work with all of our great speakers. And I'd like to thank our speakers—Rick, Roger, Sarah, and Jim—for their presentations today.

Back to you, Marc.

MARC ANDERSON: With that, I think we can end the recording and adjourn. Thank you, everybody. Enjoy the rest of ICANN.

**[END OF TRANSCRIPTION]**