
MONTREAL – GNSO - Joint CPH TechOps & REGEXT Meeting
Sunday, November 3, 2019 – 13:30 to 16:45 EDT
ICANN66 | Montréal, Canada

JAMES GALVIN:

So I have us at half past the hour. Let me just look around the room once. Is there anyone obvious that people think is missing that we should wait for?

Okay. Well, we will wait just a couple more minutes. I guess the remote people can hear me speaking, right? Yeah. So we'll just wait a couple of minutes before we get started. This is a very big place. I'll try and give people a chance to make the long walk.

Okay. We are several meetings after the half so, wow, I didn't even have to ask. Announce the meeting getting started. Thank you very much.

So I am James Galvin from Afiliac and this is our first ever joint ICANN TechOps and IETF Registration Extensions Working Group meeting. And this is a very good thing. I'm very happy that we're able to do this and to make this happen.

We won't do a round robin here and do introductions, I mean, as people can please just say your name before you speak. For those in the Zoom room, I am trying to keep a good eye there so you can raise your hand and we will try to manage speakers as people want to say things and I'll try to keep an eye on the chat.

I'm happy to take some assistance in that regard from anyone in the room if you also happen to be in the Zoom, if you notice a hand or a chat

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

and want to flag. I'm sure that Zoey will be keeping watch too and will let us know. So all of that's a good thing. All right. Next slide, please.

So our agenda for today, welcome, introductions, Note Well. The Note Well, I'm sure, is a new construct for ICANN people but we're going to take a little bit of time to talk about that when we get there. I kind of just did a welcome.

So I'm a participant a member of the TechOps group on the ICANN side but just for identification purposes, I am also the co-chair of the REGEXT Working Group. This is officially just an inter-meeting of the REGEXT Working Group. I will point out that the Zoom room has my co-chair from REGEXT also out there, Antoin Verschuren, is online and is part of the Zoom room, participating in this activity. But this from the IETF side is really just an inter-meeting that does actually have some level of formal structure and stuff that goes with it, but that, as a practical matter, doesn't really have to concern people. We'll just move on and do our work and get into it.

Let me take this opportunity then to jump to the next slide which is just highlighting what we're already doing and then we'll jump to Slide 4 which is the Note Well side and let me spend a little bit of time talking about what's there.

This is an IETF construct. It was mentioned in our TechOps meeting, so for the IETF people who are actually here, this will make perfect sense.

Under ordinary circumstances, inter-meetings don't really have to put this up there. It's really an important matter or formality for IETF

Working Groups, but given that this is being done in the ICANN context and many people probably aren't familiar with the IETF way of doing things, I thought it was important to put this up there, highlight it and call it out. It really is just an emphasis, in layman's terms, on the fact that this is a public forum to the extent that anything you say is going to find its way into an IETF document onto the publications that are available to you in that stream. It's just important to know that if you said it at the microphone and in public, then it is exactly that. It's public and it can be used by the IETF.

So if you have any reasons for not wanting something to be public, then you shouldn't speak it. That's kind of what that boils down to. There's a lot of formality around that, written up there. This Note Well has developed a great deal over time and has some very particular things to say. Another layman's point that's made in here – and ICANN has a similar kind of construct – is there is a behavior part of this. You are expected to work with other participants and be a reasoned and rational person.

And of course, even in the ICANN context, we're all very good about that so that won't be a problem even in this forum, right? We're all going to just get along fine and be buddy-buddy. But in all seriousness, there is some formality associated with all of that and the IETF is very clear about that, as is ICANN. So it's just useful to call that out.

I think I'll just pause for a brief moment and see if anyone wants to ask any questions or comment on this at all, if you have any concerns or

questions about it. And I'm not seeing any hands or looks around the room. So that's all a good thing.

And this is an open forum. I realize in an ICANN thing ... I have a peanut gallery. There's people behind me and people over here. There are still a couple of chairs open. There are no restrictions on who sits at the table and has things to say, so feel free to come on up and use the space. Okay. Next slide, please.

So with that, I thought I would spend a little bit of time talking about the objective for this meeting and why we're here and how this is going to work in part because this is the first time we've ever done such a thing, so let's talk about what it means. So let's move to the next slide please.

The objective for this meeting. It is true. The ICANN TechOps meeting is where a majority of technical representatives, registries and registrars, participate. So we have a variety of discussions about technical issues and we've been having meetings all morning on some of those issues. So this is maybe a bit of an introduction for those on the IETF side about what happens here and the way that things work.

And of course, the TechOps group has produced a number of quite valuable technical specifications for interactions between registries and registrars. And many of these really are suitable for being Internet standards and in that respect, we'd like to find a way to get a nice close, collaborative working environment here and provide a mechanism and a means to bring those documents into the Registration Extensions

Working Group and the IETF which is a working group which has two primary responsibilities.

One is to deal with technical standards related to registration services and that's, obviously, what we're going to be doing here for the most part. But it also turns out that one of our responsibilities is dealing with any issues that come up related to RDAP in particular. So the protocol that the IETF has made available to ICANN with respect to publishing registration data. So we have, essentially, those two paths for work streams in the IETF working group and obviously, TechOps has been working on a number of features and elements there and we're going to get to talk about some of those here in this context because this is where that group of people meets.

A majority of people who, on the technical side of registry and registrar interactions are in ICANN meetings. They don't make a habit of coming to IETF meetings and we haven't really had that close collaboration and close management and so we're trying to introduce that, make that happen here, and that's what we want to do.

There are as many as 14 potential documents which have originated in TechOps and we have some milestone opportunities available in the REGEXT Group. So our overarching goal here is really to look at the 14 documents that are currently sitting in a backlog, if you will, in the TechOps group and consider what is important to this group and how to organize them and structure those so that we can begin moving them through a standards process, bring them into the IETF context, get that additional external review, and have them become the technical

standards that we really need them to be so that we can point at them here in this context. So that policies in ICANN can make use of them, and in addition, the TechOps group can continue to leverage them and grow them and enhance them and that kind of thing.

So that's really where we are in this meeting. Let me take a moment here and see if anyone wants to jump in. There are several of us here in this room who really do attend and participate on both sides of this, in TechOps and in REGEXT. I want to offer the opportunity for anyone who wants to offer any other comments about this, if they have their own perspective they want to add about making this work. And I'm not seeing any hands and no eye contact so okay. Nobody wants to talk.

Okay, so let's move to the next slide here. I want to talk a little bit about process from the IETF point of view and what ultimately will happen as we come out of this group and what it means to engage with the IETF.

To a large extent, the administrative side of this really want to hide from the group in the sense that it really shouldn't be something that you need to be overly concerned about. I know that's one of the reasons why you have a Chair in TechOps, you have a Chair in REGEXT. Our goal is to try and manage all of the administrative parts of making things happen. I like to try and keep us focused on the technical work, but nonetheless, people should understand what really happens.

In the IETF, it's not actually required that you attend a physical meeting. ICANN tends to be somewhat more driven by physical meetings. The IETF has them too, but an important component about the IETF is mailing lists. Officially, all work happens on the mailing list. Decisions

are made on the mailing list and that's how you progress things along. The physical meetings are more of a means to an end. So even anything that you discuss in a physical meeting, you ultimately have to bring to the mailing list so that you ensure that it gets full exposure to everyone and everyone gets a chance to comment about it.

And that is really what will happen here. This group and the TechOps group and the fact that it produced these documents, that's all fine. IETF has a process to deal with that. It's called a design team. Design teams are more than welcome to bring proposals. We can treat the work products out of here as a design team effort and then they have to be adopted by the working group. The working group has to take them on and recognize that it fits within the charter of the group so that we can seek to make it a technical standard. So there's a broader community that has to accept that and that's what'll happen here. The goal here in this meeting is to move towards taking those 14 documents and making them into something that we can bring forward.

We don't necessarily have to finish the discussion today if it turns out that way. If we need more time, there are plenty of opportunities to meet. TechOps does meet on a regular basis. It can continue those discussions on documents that it needs to.

And then as the documents, and once they're adopted by the working group, then yes, you do have a broader community review of them and so the documents then are subject to that broader community review. So the IETF community, both the working group and the IETF community at-large, I would expect for the most part, that that won't. I

don't immediately expect anything especially substantive to come from that. But it is fair to allow for the opportunity for people to identify issues with related work in the IETF or any other kind because if you have a different set of people looking at it, they might see something that we haven't seen here and we just have to work that process a bit and deal with it.

I don't think that's any different than the way things work even in ICANN as TechOps take the domain transfer work, in particular as an example. TechOps has had a very strong majority position on something and now it's moving forward in the broader ICANN context and that's the way this will work here.

And then ultimately, upon consensus, the documents are submitted for publication and then there's a whole separate process which gets managed by the gentleman sitting to my right here, Barry Leyba, who in the IETF is the area director. He has the line management responsibility, if you will, for the Registration Extensions Working Group in the IETF. So he's fully read-in, and engaged in what's going on here and how this is working. So we at least know that we got that level of support as long as we're following process and working through technical issues, all this is going to be fine.

So that's my rather long-winded kind of introduction and engagement, and for those who were in TechOps this morning, Rick had mentioned during the RDAP Working Group meeting that it was kind of important for people to come, especially ICANN people who really aren't used to sort of this ordinary stuff, to see this and understand this.

I am open for questions and comments about it. You can do it now. You could actually, something occurs to you later and you want to raise it, that's fine too. The goal here is to make it work. That's the only objective and I will find a way to meet our needs as we go along here within the guardrails that process restricts us to.

And Rich, please go ahead.

RICH MERDINGER:

Hi. Rich Merdinger, GoDaddy. Do you have any advice or some insight you can share on folks that might want to participate in the IETF world but are sensitive about intellectual property losing the ability to come up with ideas? They want to share them in the standards, etc. but they also want to somehow, I don't want to say just glom onto it, but they want to protect it. They don't want to feel vulnerable.

JAMES GALVIN:

I guess I'll say two things. First, the IETF has a very clear policy for itself with respect to intellectual property. As a participant in the IETF, one of the things in getting into a detail that's in that Note Well, is it's pointing out that there is an expectation of behavior on the part of any participant and that is that if you are aware of any intellectual property which might have any kind of relationship, tangential or otherwise, with any work that's going on in the IETF, the IETF expects that participant to make that visible in the working group discussing that work and just make the entire working group aware of it. That is the official policy and that's the expected behavior on all of that.

I think the second thing that I would say is any concerns that you might have about that, you should take up with your own legal counsel. You're simply obligated to follow that process and what you do or don't do is up to you and your own legal counsel. Good question, though. IPR, as much as it's a popular discussion in ICANN, it can be discussion in the IETF too. Any other questions or comments?

Okay. Let's move on to the next slide then.

So now let's get to reviewing more specifically exactly what's on our agenda. In the early part of this when I had proposed this meeting, I really had pointed to people at the Best Practice Domains website. It has a very nice list of all of the documents that we really do want to look at and think about how we want to pull them together and how we want to handle that structure.

So if we go to the next slide, I have organized those 14 things into the following five bullet items that I would say are on our agenda for discussion today and the way that I'm going to suggest that we approach this is that we touch on each of these topics and have a discussion about them so that we can understand them. And know what they are and know what each one represents. And then as when we're done with that, we will then have a discussion separately as a follow-on once we're all on a page about what each of these things mean. We'll talk about our priority for these things and for those things that have particular documents, our priority for what we want to propose for adoption in the IETF so that they can progress and move forward. And we'll just, we'll see where that takes us. But I think for right

now, we have these particular five topics and we should go through them and see what they are.

I did roughly list them in the order in which I would like to touch on them and I based this order on the substantive discussion that I expected we would have to have, so more discussion at the top, less discussion at the bottom is kind of the way I had a picture in my mind about it. I could be completely wrong. It's entirely possible I am because other people may have other thoughts about what they want here. Nonetheless, this is your opportunity, the usual agenda bashing kind of thing. If you want to reorder this and do it differently, I'm open for suggestions. There are people here who are familiar with all of these documents. I'm sure that not everyone is familiar with all of them, but open for any kind of discussion and any thoughts about agenda bashing or moving things around. Otherwise, we'll just step through them in the way that they are up here.

Rick, go ahead, please.

RICK WILHELM:

Rick Wilhelm, Verisign. We can pull registry mapping from there. I think that we had talked about that previously. But we can, if it hasn't been formally pulled, we can pull that. Thanks.

JAMES GALVIN:

So thank you for that. Actually, that's why it's on the bottom of the list. I figured at best, we would touch on it, but I actually don't imagine that it's on the list of consideration for one of the slots to be adopted and I'd

figured we'd take a few minutes to talk about why that is and what's going on there but that's as much as we would say here.

In fact, why don't we just do that now and discharge it right away? So maybe I'll just put you on the spot. Do you want to talk about it a bit and what's going on there? And we can discharge that one.

ROGER CARNEY:

Thanks, Jim. So the registry mapping was, I think, two separate ideas that came together after some discussion. Verisign had actually started this many years ago for their own internal systems for maintenance and things like that. But when it came to light really was when we started onboarding hundreds of new TLDs and registries. GoDaddy produced an onboarding document of 360, 400-some questions that we asked every new TLD registry operator. So these are the questions that even if a registry is RFC compliant, every RFC has optional features in them and maybes and shoulds. So all these questions are answering all those questions that are possible and for a few other things of standardization, one registry will call something that is the exact same, completely different than another registry will call it. So that was one of the reasons GoDaddy created this and I don't know, Jim, you may know better. It's been over a year, probably two years since we've started talking about this mapping and pulling it forward to possibly solve that problem and not solve the problems for GoDaddy. But anybody that has to onboard these things, looking for a standard way to call things and answer all those questions.

So it was a discussion. Again, Rick may know the time when Verisign actually started this internally but we kind of globbed onto that and maybe made it even bigger than what Verisign had pictured it going forward. It's something we've been working slowly on for the last year and a half, two years. But yes, it's never made it onto the REGEXT dock and I don't recommend that it gets there yet. Right now, we are waiting for Verisign to solve their IPR licensing and however that works out because they do have some IPR claims based on this.

But I think, again, we'll continue working it in the background and not push for this to come on our docket. So thanks.

JAMES GALVIN:

Thank you for that. Let me just speak a little bit to clarity on the status of this work with respect to the formality of the status of the work. And I see Jim Gould's hand and then we'll go to Jim and let him add context that he wants here.

It's not, the registry mapping stuff is not officially adopted by the working group. But the chairs have been allowing the work to progress within the construct of the working group. In fact, we've had some official working group inter-meetings along the way. We've given them working group time at IETF meetings to have some opportunity to move it forward. It is important work.

The folks who are working on it have not officially asked for it to be adopted. Two reasons for that. One, it's still really in its kind of design state. It has its status and they're working through details. And two, is

the fact that there is potential for IPR associated with this and under ordinary IETF processes, we do need to understand exactly the status of that and that's just not a known quantity yet for reasons outside the control of any individual person who is part of the work. And it's just fair to make note of that.

So I mean, even the people who are progressing the details are being courteous in that they're not asking for working group adoption yet until we know what's going to happen with the IPR status and that may change as we go forward, but that's kind of formally where we are. We're allowing the work to actually happen at a slow pace, but we're waiting for that key detail to have some closure before we figure out what the next steps are within the group. And so Jim Gould, you have your hand. Go ahead, please.

JIM GOULD: Can you hear me?

JAMES GALVIN: Yes, we can. Thank you.

JIM GOULD: Okay, good. Yeah, my first note on the registry mapping is that taking a step back and thinking about registry mapping, I'm getting the sense that it's a little bit too... Aggressive is not the right word. I would say that it's very much a challenge to be able to meet or define all the various policies that are out there in the registries. So for the working

group to take that on, I think that we maybe figure out an awful lot of work that maybe very difficult for us to come to a conclusion on. So that's one reason why I wouldn't recommend bringing it in at this point.

And then I wanted to make a comment related to the data set file format. You have that under the registry/registrar reporting. It currently is not defined for that purpose. It's currently defined for defining bulk operations, but in looking at defined reports, I felt like that would be a perfect opportunity to take it and to adjust it to be able to also find reports because in essence, it provides for a formal way of defining the format of all of the various fields of a report and it's able to reuse all the existing EPP field types as well, so that's a big advantage.

So if that was to be taken on, that in essence, I would be looking for the working group to help adjust that to also support reports. That's it.

JAMES GALVIN:

Okay. Thank you. Let me just suggest that we appreciate your data set file format comment. Let's hold that until we come to talking about that particular bullet item and then we'll pick that up and dig into that. So Rich, go ahead please.

RICH MERDINGER:

Thanks. Something Jim said is very true about the complexity of all the policies and finding a way to encapsulate all of that in a way that is efficient. The sad thing is that as registries and registrars, we're already doing that today on a one-off by one-off inefficient basis. Each registrar

out there encapsulates the policies that are in play. So I appreciate not doing it now. I'm not stressing for that.

But if we want to be an efficient industry that continues to grow and registrars are able to efficiently make changes to the TLDs we support and registries get the flexibility in the features that they want to support for the TLDs, we need to come up with a way that we can. Maybe we redefine this down the road so the IPR issues for the registry mapping concept goes away, we call it something else, we implement it in a way that is a little more greenfield.

The point is the problem is there. It does exist. We all feel it. Let's address it. There's no better group than this to try to come up with a creative way to make it a more efficient industry in this regard. Thanks.

JAMES GALVIN:

Yeah, very much. Thank you for that. Let me just agree with you and acknowledge that it is a real problem and I do acknowledge that it is. And that's why we've been allowing the work to sort of continue at the moment and its role in the REGEXT Working Group. But maybe we do have to think about the priority at work and what we want to do going forward. That's important.

Rick?

RICK WILHELM:

Thanks, Jim. Yeah, I hear what you're saying, Rich, and I think to sort of expand upon what Jim Gould had mentioned, when we got into it, to

try and come to closure on the problem, the way in which it had currently been defined. It's a substantial problem to try and get everything defined, and so at times, in computer science, the road to success might be just in sort of picking a smaller problem to solve but it's not an easy problem. Interface is the type of flexibilities that registries are allowed to pursue, which is well within policy because it is a competitive market within registries and that is by design. It's not a commodity business. Registries are different. They have different policies, different purposes. There's brands. There's restricted. There's generics as we all know here.

So there are different ways in which those businesses manifest themselves, and some of that is expressed through those interfaces. So I think that one of the things, as this has been an interesting exercise in helping us all to understand that, "Wow, this is kind of a challenging problem," certainly understanding both from my own personal experience and from working on our side dealing with this problem, it's not an easy thing to solve. So understanding of the problem, but there's also limits probably to what the situation will sort of allow for given that the registries operate in a competitive marketplace and they're not going to have necessarily a plug-n-play interface. For example, in the same way in which your cell phone number moves from one operator to another, that sort of thing. Thank you.

UNIDENTIFIED MALE:

In response – and I do mean this as an “in response”, failure to participate in this, I would argue is actually anti-competitive because

we are looking at another sub-pro. The sub-pro is looking at another round of TLDs and individual registrars are not going to be able to work with registry A, B, C, D, and E, all with their unique ways of expressing their things and onboard all at the same time.

So the incumbents end up with, “Well, we’re already there. Why would be incented to make it easier for the new ones?” And I believe that is not supportive of a competitive marketplace. So I’m not picking on anybody in particular. I’m just suggesting that I’m looking at we’re a large registrar. NO question about it. We want the whole industry to thrive so that we continue to thrive. That’s what I’m looking for.

JAMES GALVIN:

Yeah, thank you. I would, and I say this with – let the transcript show that there is a smile on my face – I would really be gentle using a word like “anti-competitive” and I appreciate the fac that you’re not a lawyer. Nor am I. My degree says Engineering and Computer Science. And there you go, brother from another mother. So just like when we use those kind of terms, we need to be really gentle.

I understand your point, but we should take caution with those kind of words at a meeting like this. We have been working on this for some period of time. Gould and other members of our team here, along with Roger, and it’s not an easy problem to solve and I think that as technicians, we can all sort of acknowledge that.

And so we just need to be very gentle with that kind of language in this kind of business. Thank you.

ROGER CARNEY:

Yeah, and I guess I just want to add onto, we knew we weren't going to solve 100% of the things coming in and out and we weren't trying to. We were trying to get maybe 75%-80% of the common things that registries do do, and the other 20% that are unique, okay, we can handle that in a different way.

But instead of 400 questions, it would be nice to be down to 50 questions or something like that. And that's actually one of the goals we took into it.

And one of the other reasons this came up was the whole idea of registry transitions. It was another catalyst for why we started looking at this when an RO changes backend providers, that causes huge headaches for all registrars and could be for others. So that was another catalyst for it. Thanks.

JAMES GALVIN:

Okay. Looking around, not seeing any hands. None in the Zoom room.

So I think the action out of this is to put a pin in this. The people who are the principals in leading the discussion of this and I look at Jim Gould and Roger Carney as the two principals, are not asking for working group adoption at this time. So this is more about making sure that we have visibility in both sets of groups and everyone sees that it's there.

So I don't. I'm going to say that our action here is that this particular topic will not be part of the discussion about priority later. And so if anyone objects to that, this would be the time to speak your mind, raise your hand. I'm not seeing any disagreements so all of that's a good thing.

Okay. We are left with four bulleted items up there. Anybody want to select one in particular they want to go for? The default will be to start at the top and talk about reporting, but agenda bashing. Nobody seems to have a preference. Not seeing any hands.

Okay, so let's get to talking about the registry/registrar reporting. Before you switch slides, what I've noticed up there, what I want to just highlight is the fact that there are nine proposals currently on the BestPractices.domain website. There are nine different reports that are described there and Jim Gould was talking about a distinction that he wanted to draw about what data set file format really is versus I'm sure that he has already now caught up with the slides that I sent out not ten minutes ago because Jim is just that good, and so he's already ready to talk about how to distinguish that from the proposal that we're about to jump into and talk about. And that's fine.

If the documents can work together, then I think that that's awesome and we can sort of work through that and talk about it here.

I'll note as a matter of formality, the data said file format document is actually, at least in an IETF context, officially expired. But that's just a detail. Quite honestly, it's easy to fix that. Not a problem. So I just want to call out that I'm aware of that in case anybody else notices and wants

to mention it. It's an easy fix. It's not a problem. We can still talk about the substance there and how it relates to everything else that we're doing and just treat it as a live document and proposal.

And with that, let me ask that we switch to the reporting slides.

UNIDENTIFIED MALE: The answer is no.

JAMES GALVIN: It's okay.

UNIDENTIFIED MALE: You can always ask.

JAMES GALVIN: That's alright. Patience is a virtue here. You've got to give technology a chance to catch up.

So I'll point out that this is a proposal that is coming from Afilias, from myself and Joseph Yee who is in the Zoom room and online with us here. I will walk through this and Joseph is certainly available to talk about this and represent it.

This is a set of slides that we've derived from a document and some work that we've been doing in thinking about all this. And I'm hoping to get a lot of good discussion here. We'll see how all of that goes. And then

our plan, actually, is to have an Internet draft that we can put together and distribute prior to the next IETF meeting.

The next IETF meeting is in two weeks and so we've already put in to get an "any other business" slot for new work on the IETF BridgeX meeting that'll happen in two weeks. So we'll have some discussion there and we're hopeful we'll have a positive response here and then with all of that, we'll be able to seek some working group adoptions. So we'll see how the discussion here goes.

Okay, next slide, please.

So on slide two, this is really intended. It's sort of your very simple text-based look at what a report currently looks like and it really is just three examples and it shows them up there. There's a transaction report, a premium name report, domain info report. The transaction report is one of those definitions which is in those nine documents which is there in the list of best practice domains. There's a premium name report. Afilias happens to have one, which looks a lot like what's in the TechOps document, and so we've just kind of listed that up there.

So the list of fields that are in that report, the column headings that go with it. And then we picked a report which was ours and the way that we happened to do it, just as yet another example to show things. But this is just a static example of this is something which exists today and reports that are present and most registries produce something like these things available to registrars. So I just wanted to put that out there. Okay, the next slide.

So the question is: What are we trying to achieve here? What is our goal? What are we trying to improve and how do we want to make things better? Right now, what's out there is this notion of let's have nine different documents which describe the reports and every time we want a report, we'll produce a new document that sort of says, "This is what we want." And when I think about, well, what would be a different way to do that? The situation today is reports are static. The registry gets to choose everything, right? They choose the format, the content, the syntax.

Different registries do different things and so I know that the registrars in the room are quite familiar with this and they have to deal with this problem all the time, so it would be very nice to be able to do things a little bit differently and make this look better.

Obviously, that was a principal motivation in the set of nine documents that were produced. And in today's world, the registrars are, frankly, in a place where they just have to comply with whatever the registries have done.

And to a large extent, this probably worked okay prior to the 2012 round of new gTLDs. There weren't that many registries and there were a lot more registrars. And so that probably was a model that was workable. Clearly going forward where there were a lot more registry service providers and even more registrars as we go along here, we need a different kind of model that just doesn't scale to have all of these differences.

So the place we want to get to is it'd be nice if reports were still static in the sense that we know what they are and everybody can do the same thing, right? That's really what we want to get to, is kind of a standard set of reports and then a standard way in which to define future reports so that they're fairly easy to create and it doesn't require a huge amount of work on the part of anybody, whether it's a registrar or a registry. We roughly know what a report is going to look like and you simply build out the little edges that you need to, to make it work.

So you get to a place where the format of the report is standardized, the syntax of the element is standardized, and the reporting elements are standardized. You know what's available to you in terms of what you can put in a report.

And then you generate reports with standardized elements. Those reports themselves are also standardized. There is, in general, there is a basic set of reports which is fairly common and we should make them common. The idea here is as an opportunistic thing, there's a base set of reports that all registries could produce and produce in the same way for all registrars. And then along the way, we provide a method for easily creating additional standardized reports if we need them or even providing some optional elements in some of those standardized reports if a particular registry has additional data that it needs to put out there that's relevant. You provide a nice, easy way to make all of that happen, again, in a standardized way. So it scales nicely and conveniently.

Those are the goals and the objective in what we are trying to achieve here. And then we got to thinking about what is a straight-forward way to do this and so this is our look at how to do that. So next slide, please.

The model here is for those in the IETF are familiar with this, and hopefully all the registrars in the room are familiar with this anyway, there was an EPP Extensions Working Group which was the predecessor to REGEXT. And what it did at the time, was to create a registry with IANA of EPP extensions. And that is the model on which we based what we're proposing here or what we're doing.

That EPP Extensions registry is really quite simple. It's a place for if you have an extension, an EPP, that you use. You believe that it is a good way, if not the best way to do something, that you can then document that, whatever it is, and stick it out there and then everyone can point at that. The idea is the market will simply find a way to say, "Oh, I want that. Let me go ask people to do it this way." Registries will do it that way. Registrars will ask for it to be done that way and so we create a model in which there's standardization in this space.

I'm sure that many registrars are quite familiar with the fact that there are a lot of EPP extensions. There are only a couple dozen right now listed in the EPP Extensions registry, but it's there and hopefully over time here, we can grow that.

So based on using that model and going forward, the specific suggestion is twofold. One, create a registry of column headings. So you create a column heading registry which has the information that's listed up there. There's an ordinal value for the entry and you'll see why

that's important in a moment here. But every entry is numbered so for every column heading that exists there, there's number one, number two, number three, number four, and you'll see an example going forward that everything gets a number.

There is, obviously, a name given to the column. Whatever the date is going to be, so there's a standardized name for that column. We agree upfront what that's likely to be. Sort of the obvious things is as you go back and you look at the example reports, there tends to be a standardized heading for a lot of things, but not always. And so let's make them all the same so we know what they are.

Obviously, in the IANA registry, you need a reference to the definition and you need to know who is registering it. So in the case of anything which gets built after our initial pre-population, if you believe that this mechanism will serve the purpose that we need, we will pre-populate these tables with a lot of information, a lot of column headings and reports, and then you'll just be able to use them. And if you need to add anything, you can.

And there's a few other administrative elements that are out there that I reference for completeness but I actually won't be talking about them directly here. If you go look at the EPP Extensions table, you'll see some of the other elements. But that's really more of an administrative matter as part of setting things up properly in IANA. It doesn't affect the substance of the technical idea here.

Now the next thing about this registry to understand is the EPP Extensions registry is what's called a first-come, first-served registry

and that's what we're proposing this registry would be also. A registry of column headings. And first come, first serve, registry is a defined construct in the IETF. It's a well-defined construct. There's a number of opportunities for what to make a registry in IANA. We pick this one, the idea being that assignments in the registry can be made by anyone, literally on a first-come, first-served basis.

So we will pre-populate with whatever we believe is a good standard set to get us going and then the future is anybody who has something else they want to add, you just do it and there's no restrictions. There is a review that will happen on any proposal that comes to IANA for a new column heading but the purpose of that review is just to make sure that it doesn't duplicate something that's there because you don't want duplicate entries. The whole point of standardizing is to have one of them. And that the request itself has to be well formed.

So when you create a publication of the new column heading, you create a definition of it and that has to be published in some way that's easily referenced in this table going forward. Then they just want to make sure they've got a well-formed reference and as part of developing the document, we'll have to specify what it means for it to be well-formed.

In the EPP Extensions registry, part of the RFC that was produced that created that registry, it actually has a whole section which talks about what it means to be well-formed. And there's a group of people who do that analysis on behalf of IANA And those are details that will be worked out but that's all that's going on there. I saw a hand. Roger?

ROGER CARNEY:

Yeah. Thanks, Jim. So speaking of the EPP extensions, there is an expert group that reviews any new extension. Again, just to your point, looking for duplicates, looking to see if there's consistency issues and if it looks good.

I think one of the keys for if we go forward with something like this is the reference. Does that reference really tell you what that is? Because that's probably one of the biggest things when registrars are looking at different registries is create they could have 16 different types of names, crdate, datecr, and you've got to find out what that really is. And I think for registrars to get use out of this, the references would have to be pretty specific so that it is. So I don't know if it's an expert group that has to do that. I don't think IANA would take on that chore of confirming that that reference was valid or made sense. So it's just something to think about. Thanks.

JAMES GALVIN:

So let me expand a little bit in response to that. I'm going to expand a little bit on that last bullet up there. There's no substantive review of any request for additional things. But it is checked to see if it's well0formed or that it's not a duplicate.

The way that it's done with the EPP extensions registry is there is an identified list of people. Scott Hollenbeck from Verisign is actually the primary on that and I think there's six other people who are ont hat list

and I think most of us are in the room here, in fact. I mean, I'm also on that list but there were some others who were on it too.

And what happens is IANA, those people are identified as part of IANA's management of the registry itself and that can change over time. And what IANA does when request comes in, because they manage the whole process here, is they reach out to that group of people and in particular, they reach out to Scott as the primary and then he solicits comments and looks for somebody to review it from his expert team.

But in addition, the REGEXT mailing list currently is part of the way that we do that. We actually use that mailing list as a way to expose the fact that a new request has come in, it's to be reviewed, and this has actually worked. There are some worked examples of this in the REGEXT group. So Scott doesn't do this in a vacuum, just among those small group of people. He keeps the mailing list engaged in a while and I would expect to do a similar kind of thing here.

And in fact, maybe a detail to talk about is maybe the TechOps group is sort of the right place to make sure that's visible too.

But all that is specified in the RFC. So I just want to be clear about that. It's always exposed to everybody. It becomes visible. But there is an identified list of people who make those two choices. They inform IANA about whether it's well-formed and if there's no belief that it's a duplicate.

I'll get to the syntax one that you raised in the next slide when get there, but did I answer that part of your question? Anyone else? Rick, go ahead.

RICK WILHELM:

Thanks, Roger. Thanks, Jim. I think I heard maybe a slightly, maybe a deeper part of the question that maybe you didn't get quite as much answer. It was around maybe the detail around almost the middle sub-bullet there. It says reference to the definition and I think I heard Roger asking more about what qualifies as the definition of the particular field because the semantic, the meaning behind the particular field, is oft times the sticking point around when it comes to interpreting a report in whatever context it is, whether it be a date or a value or a total or something like that and how are those going to be? That's what I more pulled out of the question and something that, frankly, is a question that I have around the whole effort, not just around this proposal, but a key question around the whole effort. Thank you.

JAMES GALVIN:

So let me use that as a "That's what I meant by the syntax question". Maybe I was a little overly obtuse about all of that, so thank you for calling out the detail again. Let's move to the next slide and let me speak to that issue a bit with this particular example.

So this is, in essence, what the column heading registry would look like. Okay. Essentially, four columns of substance. There is the ordinal reference number so that would be part of it. There is the name for the

column heading and we sort of listed some. What I actually did was pick the domain info report and just sort of listed those column headings in there for the moment.

In this particular case because this would end up being a pre-populated entry. There would be an RFC and whatever RFC we produce here which creates the registry, as part of that, just as we did with EPP extensions, you would also specify the pre-populated values and what they are.

The registrant in that case would then be the ISG because it's coming out of a standard as a work product of the working group. In the future, if someone had a particular thing they wanted to add, then they would be listed as the registrant. And that's the way the EPP extensions registry also works.

The first dozen or so all have ISG as a registrant because they came out of the working group and a work product. So in particular, what I would expect needs to happen in this, if you accept this sort of model about this, a part of, there are two essential parts of this document to be produced. One is about creating this registry of these column heading names, and so in the document, each one of these column heading names would have a section which would give you the name of the column as it's decided and then it would tell you everything about that column and that data which is relevant, so the syntax of it. In particular, if it's a date, exactly what the date syntax is supposed to be.

If it's a status value, what are those statuses and where do they come from? And what those values are?

The restaurant ID, where it's defined. In that particular case, that's something which would come out of the EPP spec so you have a reference to it in there.

The reference here is intended to be the document which tells you how to find that information. So the details of the column heading have to go in the definition of the column heading.

And so it becomes a standard and everybody has to do it the same way and that's what it becomes. We specify exactly what it looks like. So if you're going to use that column heading, then you're required to stick the data in that report in that way, however, it's defined.

And this would be the discussion that we have. Sort of the fairly obvious thing to do is you take all of the report definitions that are proposed, you just gather up all the column headings that are there, and then you work through them and you collapse those that are similar. You reach out to the community for any others that are not there that maybe we want included.

Part of this document also has to specify what it means to be a well-formed entry on this list. And so that's also what I mean by the syntax. I just, the obvious things are if you have a date, you want to specify exactly what the date field is supposed to look like, and that will be in here and that's part of all of this too, and it would be there. But you have to define what it means to be well-formed and what it means to be well-formed is that you would specify what all the syntax is for anything that you put there and that's what gets looked for.

I think that's all I want to say about that for myself as an introduction.
Jim Gould, do you have your hand up? Go ahead, please.

JIM GOULD:

Yes. This is Jim Gould from Verisign. I noticed that there's no concept of name space in this. Is there a need for that? Like for example, status. Are you talking about domain status, host status, contact status, which have different enumerated values?

And also, I have a question related to the format definition. Are you intending to use [inaudible] schema, [inaudible] schema or some formal way of defining the accepted format for that particular column heading? Thank you.

JAMES GALVIN:

So thank you. Yes, with respect to the status question in particular. One of the things that we will have to examine is, and I chose status and I put it up there just like that on purpose, status is a column heading that's used in a number of reports and I don't believe it's actually the same in all reports. The syntax is actually different.

So just as you said, we would probably have to put an adjective in front of the word "status" and create different column headings for those different elements that have to be there so that we can properly define each one. So that's one piece of what you said.

The other piece of what you said is in terms of whether this is an XML or something else, that's open to the group. Well, in terms of what's in this

document, I expect it just to be a textual statement. This is not a machine parsable document, this standard, so I wasn't expecting that if we were to define status, that we would have some kind of XML representation of what that is.

We would simply state what it is and draw from how it's defined in the context in which it's used. So when you're talking about status from the point of view of locks or status from the point of view of whether the name is active or in a renew state or a delete state or things like that, depending on where the status is used, then we will refer to how those things are already defined in the IETF.

And the same thing with domain name. I would expect we will refer to the definition of a fully qualified domain name and we'll have to make a decision about whether it's an A-label or a U-label, those kinds of things. Those are the details which we will come to and they will be all part of this in detail, and the same thing with the date.

Did you want to respond, Jim? Go ahead.

JIM GOULD:

Yeah. The question was does the column heading need to be unique? So in essence, would the column headings be defined based on report definition which refers these ordinal reference numbers to identify the meaning behind a referenced column heading, like status?

JAMES GALVIN:

So I apologize. You did ask about the name space and I apologize for not responding to that.

Yes. The part of not being a duplicate entry is that you're not allowed to reuse a name that's already in use. So the choice for the name can be anything that's relevant to you unless we decide we want to restrict it. There's always sort of the obvious restriction of U.S. ASCII for labels but I didn't want to prescribe that at the moment. I wanted to allow an opportunity for some thought to be given to that. So it is true that the column headings do need to be unique. That certainly is a requirement. And the label can be anything that the definition wants to make it, and then you just define what it is. So please go ahead, Neil.

NEIL MCPHERSON:

I'm Neil McPherson from [inaudible] for the record. I wonder if we don't need to take a step back and think about who they are designing this for, right? So it should be that registrars can simply consume these reports and there are simple ways of doing that, right? We could just write it down on a piece of paper to say, "Hey, these reports are looking good," and the registries can commit to taking a bit of paper and using it.

And then every registrar knows that they're getting the standardized report from the registries or there is something like this, which is to me seems kind of over-engineered for what we're trying to achieve. I would love to hear, and I think [inaudible] with regard to the reports, there's kind of like three, four or five reports that all registrars are using on a daily, weekly, monthly basis that maybe need kind of firm

standardization of. There's others that are maybe kind of extensionable or we don't know what the future brings. Maybe we need some other stuff which we can kind of work on, on the fly. But the major use cases for transaction report, transfer report, blah, blah, blah – four or five things that we all need.

So I'd like to hear why doing it like this is an advantage to just writing an IETF best practice draft and everyone working and using that.

JAMES GALVIN:

So let me use that as a segue into, if you don't mind, to jumping into the next slide, please, the registry of reports and talk more about how all of this is supposed to work.

So again, as I had started with the registry of column headings, the idea here is to create an IANA registry of the reports to be produced and it'll be modeled also after the EPP extensions registry, and of course now, after the column heading registry. So the complete proposal is to have two registries at IANA, two IANA registries, one which is of the column headings that are used, and then the second one which is "Which column headings are to appearing which report?"

So we make a registry of the reports that are being created. So in fact, if we would just do a one-to-one mapping here, what I would imagine is we take the nine documents, fi you will, in a very gross level and simple case. You collapse all the column headings together and you get them all into the column heading registry. And then this particular registry would have nine entries in it, one for each of the reports. And the

registry would have in it the elements which are listed up there. So the first one, of course, would be a report name. You'll see why that's important in a moment.

There needs to be, again, a reference to the definition as before so you just need an RFC number or some other publication. There needs to be some other kind of URL that references the publication for future things. And the important element here is the ordered list of column headings in the report. So whatever column headings you want in the report, you select out that ordinal number and that's what goes here. And you put them there and you list them in the order in which you want them to appear. And that's what standardizes the reports. Okay. So, if you want a particular report created, that's how you do it.

The other thing is the registrant to the element and a few other administrative elements that go with it. Again, this is a first come, first serve registry.

Let me jump to the next slide and show you what an example looks like. And I only put one in there at the moment because I didn't actually ... I was trying to figure out how to make this thing look like something you can display, and if you try to put them all in there, the font gets too small and nobody can read it and you can't see what's there. But hopefully this gets across the point here.

We have the Domain Info Report. So, we give it a particular name. I took that one. It has a reference to an RFC, which is whatever one we create here in this sense. You can see the column headings there. I deliberately only moved one of them – the TLD one – to the end and I left the report

as it's currently defined. So, the columns are 7, 1, 2, 3, 4, 5, 6. That's the way that report is defined.

So, now what happens is when a registry says, "I want to have this particular report," you simply look at a registry and you say, "This is the report I want and you should produce it this way." And if there's a new report that a registry wants to create or that a registrar wants, there is again a first come, first serve mechanism for creating that report, specifying it here, and then you know that everybody will do it the same way and you just have to deal with your business arrangements to make something come into existence the way that you want it. So, let me pause there. Go ahead, Jody.

JODY KOLKER: I'm curious on the reference there. So, it's just a domain name full report. I shouldn't say "just" but it's something that we already have, basically. Do we need to have an RFC for it, then?

JAMES GALVIN: The RFC number here ... Yes. There needs to be a reference to the fact that that's where the table entry comes from. So, it's a well-formed request. You have to define the report, and then all you're doing in the table at IANA is capturing what's in the report.

UNIDENTIFIED MALE: Roger was explaining this to me. It's not a separate RFC. We don't have to have another RFC that explains the domain name full report and then

puts it in there. I mean, this could be the RFC to 5713 or whatever it is for the domain, for EPP for domains, right?

JAMES GALVIN:

This could all be one RFC. Fundamentally, there's actually sort of one obvious alternative. One way is in which we do all of this together into one giant RFC and that could be the way that you do it. It's sort of an obvious alternative is you could have two documents because you could have one which is defining the column headings and then a separate one in which we list the nine reports and make that specification. It would be reasonable to do that. That's sort of an obvious alternative because they're two different registries and this way you can talk about them separately. But there could be one giant document. Or we could also do entirely different things.

I think my main goal in this proposal was to get away from nine documents and then set up a mechanism by which other things can be done. And we can certainly pre-populate with whatever we think we need. Did you want to raise your hand? Okay, Neil.

NEIL MCPHERSON:

Maybe I'm a little sensitive because I wrote a couple of those documents but it seems like [inaudible] nine documents is kind of like a dirty word that keeps coming up. So, the number nine is also going to back to what I was saying. I think most registrars or registries don't even touch three, four, five of those. I would say 90% of the registrars only need domain transaction reports. They're only going to be dealing with

one of those current drafts. So, we should get away from the fact of ... We shouldn't be thinking about trying to fix the problem of we've got nine drafts out there, trying to fix the 80/20 – fix 80% of the problem first and we can probably do that with two or three of those drafts currently and then other 20% is the other six or seven.

JAMES GALVIN:

Yeah. That's perfect from my point of view. If we decide that we only want to focus on a couple as the baseline that you really would like everybody to do, and then everything else just becomes optional to be dealt with later, I'm good with that. This is more about the mechanism for causing these things to come into existence in a standardized way. It's sending up a standardized path for this kind of stuff to happen in a relatively straightforward way.

UNIDENTIFIED MALE:

So, I just kind of want to walk through this just a little bit because I think what I'm sensing and what I have in my own mind, too, is that as soon as we say RFC, that's when everybody's eyes start to glaze over and we don't get anything done. This would have to be ... Basically, we'd have to have an RFC that would have to explain the headers in order to get the registry out there. Is that right? We would have to have that done first before we could actually have a registry.

JAMES GALVIN:

Well, yes. No, all of this can be done together at the same time. We're going to create one document and pre-populate everything. But, yes, in

order to get a report, you have to define the column headings and then you define the report and that's it. And we will put both of those things in this document. We will pre-populate these tables, these registries, with the baseline that we're looking for.

UNIDENTIFIED MALE: Sorry. And we can't do that without an RFC?

JAMES GALVIN: Well, it's fine that we keep saying RFC. You're right. And that's because in talking about it in terms of the way the IETF understands it, the IETF publication series is an RFC. So, think about this as anything you want. It is a persistent archived work product that will have the definition in it that can be referenced in perpetuity kind of thing. That's the deal.

UNIDENTIFIED MALE: Okay. So, then, every new column header that goes out there is going to have to be part of an RFC, no?

JAMES GALVIN: Actually, no. This one will be because, as part of creating the registry – creating the registry has to be done as an RFC in the IETF context. So, doing those two registries really does have to be documented that way. There's no other way for the ISG to ask IANA to go do something unless it's documented that way.

But what we will do, as for our own purposes, is we will pre-populate the tables and include that in this document. And then, in the future, it's a first come, first serve registry. So, what we're doing in this document in defining these two registries, we're opening the door for anyone to come forward with any other column heading that they want, any other report that they want – literally anyone.

All they have to do is put up the definition of it, publish it in some way, which can even be just a URL to a document on your own website. It does not have to be an RFC – unless we decide we want to make that constraint to it. We could make that constraint to the table but the IANA requirements won't require that unless we say that as part of defining the table. We have to decide what well-formed is. That's up to us. But I'm guessing that we won't make well-formed be that it has to be an RFC, so that if somebody wants to invent something, especially if it's an optional thing or has limited scope and limited applicability, you just want to make sure everybody knows how to do it for when it's needed, then you simply put a document up. You ask IANA to make the entry. You go directly to IANA to do it. You don't have to bring it anywhere special. And then IANA works a process to get it reviewed and then it gets an entry and you're done. And then you just get to refer to it as needed.

UNIDENTIFIED MALE:

So, yes, Jody. You could do this outside of an RFC and it would not be with IANA. So, we would have to create a registry that would hold these values and then create something that refers people to that registry.

So, if we use IANA, it's already established. And if you look at the form here, there's one or two RFC numbers that are going to go into this for all of them. It will be the original one or two that creates the registries. And that's the RFCs that are going to be referred to in here.

UNIDENTIFIED MALE:

So, my question is – or what I think that we're trying to get to, and this might be for a sidebar conversation, is that we already have that. It's called BestPractice.domains. But nobody is paying attention to it. So, we already have a URL. We already have all that. We already have this out there. But in order to get anyone to listen to it, it seems like it's got to go through an RFC, right?

JAMES GALVIN:

So, I will respond directly to that. I know that you have your hand up there. So, actually, no. It turns out one of the paths that we could take here ... This is a working meeting, so let's look at our options and get them all on the table. You are correct. In terms of, for the report registry, the reference in the report registry could be to the specification that is on the BestPractices.domains website. So, we don't actually have to bring that inside this RFC. But the column headings here would have to refer back to column headings that are defined in the other registry.

And it's not immediately obvious to me that we could escape having those in the RFC, quite honestly. They probably want to have all those standard column headings in a document in that way. But we can

continue to talk about that if somebody has a different opinion about it, as far as that's concerned. Go ahead, please.

ANTHONY EDEN:

Anthony Eden from DNSimple. So, just using the EPP extensions example, their URL is in there in cases where there's no RFC. I will point out there is at least one case where the URL no longer works, which is a consideration, because if something shows as active in the registry, there needs to be something that covers the case where they break their URLs and the advantage of going through an RFC is ultimately very permanent.

But I agree with you 100% that it should be defined is, at minimum, a URL somewhere that is functioning, for the reference.

JAMES GALVIN:

So, my own personal bias in this – I am pushing this forward and trying to make this work in an IETF kind of way. My personal bias on that is, because the IETF represents a persistent, permanent historical archive. So, if you get it published through that cycle, then you don't have those kinds of problems.

I also acknowledge that it's really only appropriate for stuff which really is a standard. Everybody is doing this thing and you really should all do it the same way and that simplifies everybody's world and you put it there and it's a done deal.

I imagine that, for anybody who does an optional individual entry into these things, you have your own obligation to maintain that. And if you fail to maintain it, then it's just not going to be implemented. I mean, if you ask somebody to do it and you can't maintain a persistent reference to it, that's on you. Nobody is going to want it and nobody is going to care. But that's sort of the downside of a first come, first serve registry.

But I still think first come, first serve is the right thing to do. The IETF does have registries that have very strict rules. You can create registries with really strict rules if you want. I just, at least me, I don't think that that's necessary. Others may feel differently. If you want to go into a stricter thing where it has to have a technical directorate that reviews it, there's an authority that gets to decide whether you can do it or not, we create a registry like that if you want, if people think that that's more valuable. I just, at least for my bias, it's not my first going-in position. Go ahead.

TOM KELLER:

I have to admit I'm not the most savvy person when it comes to IETF processes. Can you give us some kind of a time scale this would be on? Is it a one-year endeavor, a half a year, is it something that depends on how people participate? I really have no picture. I mean, I don't mind doing it that way, but if I'm correct in my assumption, this is just basically using a different format and it's not binding on registries either. So, I want to ask the question why are we doing that, really, but that we have a well-documented reference.

JAMES GALVIN:

Your binding question is an important one. And let me just speak to that directly. Nothing that we're doing in TechOps, and neither is anything we're doing in the IETF binding on anybody, so if you're looking to say you have to do this, this does not achieve that in any way and neither does having it on best practice site, or at least in ICANN at the moment. From an ICANN context, if you don't have a consensus policy that says this is what you have to do, you're right. There's no binding at the moment.

What you get from this is a persistent, static definition, which really – I mean, I said this in prior discussions that we've had about the BestPractice.domains website. I like it. I think it's a really good thing and it's useful construct and all of that. But to be honest, it does not meet the requirement of a persistent archival publication series in any way because it's just at risk. At least, in the IETF context, if you get it that way, then you have that.

If you're looking for something that you want to be standardized, because you want everybody to do it this way, there's still the binding question. But if you want to make sure you've got one definition that will survive for long term, then this is a path. I'm certainly open to other suggestions. I mean, there are other standards bodies. This one just seems fairly obvious and an easy one to get to. So, that's a critical criteria as far as that's concerned.

I think there was another point that I wanted to make and I feel like you asked a couple of ... Did I respond to everything you asked or did I forget something?

TOM KELLER:

No. The first thing I asked about was the timeframe. So, how long will this whole endeavor take?

JAMES GALVIN:

Timeframes. I believe – I sat down once and I did this. You might know offhand. On a happy path, if everything works on a happy path the way that you want, you could actually get an RFC out in something like three or four months, something like that if you actually ... Because just like ICANN, you have windows for how long it takes to do something because you have to public comment periods. You've got to have last calls. You've got to do all of that.

Honestly, as a practical matter, I view this as a one-year effort. I'm really looking for whatever we get done here, this will be out the door before the end of 2020, somewhere along the way. And what does that depend on? If we get together and we actually ... Whatever amount of time it takes to develop the specification in detail, cover everything so that we like it, so that's a variable amount of time. We're sort of talking about the issues here and the concept, but ultimately, this all has to be written down and carefully reviewed and everybody has to have a say. Then you start the process side of it, which matters. The working group can take

this on. You can progress the work as quickly as you are willing to sit down and work the document.

Once that happens, the next steps are working group last call, IETF last call, and then you submit it to the RFC editor. Those things require that you be on the mailing list and actually submit your votes and make your comment known when the time comes. But the work can progress as fast as you want.

Publication cycle. A week for IETF, where working group last call, IETF last call is two weeks as long as it comes out of a working group – or it's four weeks. I'll let Barry say what this is.

BARRY LEIBA:

Best case, when the working group is done with it and sends it to me, best case is four weeks. I would say six weeks is a good round thing. Then it goes into the RFC editor queue and that just depends on the workload the RFC editor has. And right now they have a pretty big workload so that can be a couple of months at that point. So, figure three months from when the working group is done with it to when an RFC gets published. And then it's however long the working group takes with the document from there.

JAMES GALVIN:

I really do think that, as a practical matter, as long as we are committed to it – and we're pretty well done is part of the problem. I mean, not the problem. That's the feature. We have the documents. It is my intent to just copy and paste, put it all together. We got to wrap up the definition

of the registry in it. The working group doesn't have to take a lot of time here. But it really is under our own control. When we're done, we're done, and then the rest is just process and it's up to the chairs to make sure that they move things along. Barry has his part of the process and you will move things along.

UNIDENTIFIED MALE:

Experience says that the vast majority of delays in getting documents done are the document editors, the document authors, themselves not responding quickly. If the authors make quick updates and push the working group to review and comment, you can get things done very quickly.

JAMS GALVIN:

I know you have your hand, Rick. Just to touch on that for the ICANN people who are not expressly familiar with IETF processes, the important thing to recognize is that the authors of the document who have ownership of the details have a role in every major step along the way. So, there's the working group process, there's the ISG review process, and there is the RFC editor publication process. The authors of the document, the named people in the document, have actions that they have to respond to each one of those steps. And if you don't keep up, then that causes delay and that kind of thing. Rick?

RIK WILHELM:

Thanks, Jim. Rick Wilhelm, Verisign. I'm not going to talk about process. Sorry. Sort of getting a little bit more back to the scope and design

target. I think, in looking at this and such, it's my ... I think I'm pulling away from it that it's my understanding that it's not really a design goal of this solution to try to enable automation based on this. It's more meant to enable the exposition and memorialization of report definition, which is quite a mouthful and I apologize for that. But automation is not a design goal. And I'm not casting anything around that. I just want to pressure test my understanding of what the objective is of what this does because I think there's a lot of value to doing that, to memorializing and making permanent and externalizing the report definition and formalizing it. That's valuable because ...

Roger brought up a good point about the precise definitions around the semantics of fields and just starting to [disambiguate] but it's my understanding, based on what I've heard, that automation is not necessarily a design goal. And that's not a loaded question. I'm just, like I said, pressure testing my understanding. Thank you.

JAMES GALVIN:

So, let me respond to that by asking to jump to the next slide. And let's be very clear about what we mean by automation and what it looks like. You are correct. This specification here is not intended for automated consumption. The RFC is intended for the developers to consume it because they will write code to consume the reports.

So, the idea here is to specify what the reports will look like, so they can be consumed in an automated way. So, that's where the automation is and what you get.

And more importantly, the automation is standardized. That really is the key goal here. Laying that out in a persistent archival way so that we can then build whatever else we want around it. This gets back to the binding question that Thomas was asking.

Certainly, it is our intent – I will speak for us and I think there are others here in the room who would say the same thing. Whatever is decided here, we will implement and we will move towards.

But, the community, if we want this, we all have to agree to move towards it. Even without it being binding, if we're not going to move towards it, then we actually do need to figure out how to solve the binding question. But that's a TechOps question that was discussed this morning, too. There's a question of once you have a specification for how you want something done, how do you make it the way that it is done? How do you cause that to come into existence?

The market will have some pressure. Not to call out GoDaddy but as a significant player, I imagine that once all this stuff is done, they're going to start looking at registries and say, "Go do this. I want it this way." Odds are, many registries – if not all or most – will align with that because if you want GoDaddy on your side, that's what it takes kind of thing. That's just a practical reality. I don't mean to pick on them as trying to force anything. And I would expect that some of that will happen with all of us in the way that we do this.

So, in terms of automation, the idea here is part of this specification in terms of the report, we'll talk about reports are created as CSV file. That really is the model and the intent here. So, everybody will do it that way.

The first line of that CSV file will always be the standardized column headings. The rest of the file will be the data that's there for that column heading. Unrecognized column headings will be ignored. And it would only be added on the end because the report registry tells you which column headings and in which order and that's the way you produce the report. If, for some reason, a registry wants to add some [optional] data there for whatever reason or in whatever context, it doesn't matter, they would put those at the end so that others who get the report who don't care about that can just ignore those column headings.

In particular, if you happen to unfortunately use and [inaudible] non-standardized column heading for something, the idea is you ignore it and that's the way that that works.

Then, there's a proposal here for a file name specification. It's a fairly straightforward and obvious kind of thing. Certainly, there are other alternatives. This just seemed like the obvious idea.

The main reason for a file name having a standard name is it makes it human readable. From a certain point of view, the machine doesn't really care if you're going to consume these reports in an automated way. But you want to give it a machine readable form, so when you're doing a listing in a folder or a directory for these things, as an individual, you can fairly easily pick out what you want.

But the initial going in position here is the date, TLD name, the name of the report. This is where the name of the report is important in the name registry, in the report name registry. That's where that comes

from and this is its use. It gets stuck here. Then there's a version number just in case you happen to produce the report more than once on a given day.

There's some other stuff that goes on after this but let me cause here and get a question. Go ahead.

BEN: Just a minor thing. Could we suggest TLD name, date version? Because it seems more intuitive to me that TLDs, like their files be grouped together rather than everything from a particular date being grouped together. Also, data and version are very similar. Date is effectively the major version and version would be a minor version within that date. No? That's kind of how I'm reading this.

UNIDENTIFIED MALE: Real quickly. I could envision a rerun of a report that is dated later but it's actually ... I don't know. I wouldn't want it ... It depends on how we define date, doesn't it?

JAMES GALVIN: Yeah. Let me phrase it this way. I think these are the essential elements that need to go in a file name. We can certainly have more discussion about the order of those elements in the file name and the relationship to each other, which I guess is what Rich was just speaking to. Is version really part of the date or is it a separate thing? That's open for discussion. We were going to propose it this way in detail and also make

them independent elements and then just consider that a baseline and initial proposal. If folks want to think about ...

What's important here is think about how your environment works. Think about what's going to work best for you. Then let's bring that to the discussion. This kind of stuff is a detail. From an implementer's point of view, for me, I don't really care what order those things come in. There's some level of open flexibility for what the elements mean to each other. Happy to have some more discussion about that. Rick?

RICK WILHELM:

A key detail for implementers everywhere is coffee. It just so happens that at 15:00, there is a break and I'm just ... In two minutes. So, I don't think that we should truncate the discussion. Certainly, I think that we should bring back. And you're the chair. I had gotten an alarm on my thing saying, "Ding, ding, ding!" So, thank you.

JAMES GALVIN:

We got two minutes. Tell you what ... Go ahead, Neil. There are two ore slides which I think we can just touch on quickly and still get to break at 3:00. Go ahead, Neil.

NEIL MCPHERSON:

Yeah. I think that the memorialization of the standards is the key, right? And we've got the BestPractice.domains thing on one side. We've got this proposal here. Then I'd love to talk about some kind of middle ground as to how we could get there maybe on a faster, easier track,

because I'm still not that convinced that having this memorialization and whatever, in whatever forum, is going to be enough for registries to implement anyway.

JAMES GALVIN:

Okay, excellent. We'll make that as a topic to come back to when we come back from break. Also, Jim Gould had raised the question earlier about the relationship with data set file format. I want to give him a chance to speak to that question.

Let me quickly jump to the next two slides just so that the people have all the information. An open question that I don't really have an answer to is: do we need to solve the problem of files so large they should be split? This is different than version number. This is about a report which is just so big it needs to be in separate files.

There's the obvious way of doing that. I chose sort of an obvious engineering way to represent the file name where you stick a [inaudible] number in there and you do that. This is just there for discussion. I don't know if we need to solve this problem or not. And it really is a question to the community. As you, as a registrar or you as a registry, if you produce large things, what's your relationship? Anybody have any experience in this and do you think this is a problem that we need to solve?

Let me move to the next slide, then, too. Part of this in all of this is where this is about creating these reports and how you deal with them. The opening proposal here – and this is subject to discussion because this

is one of the areas where the data set file format document and its proposal actually covers this because it gives you a great deal of flexibility in this space.

Our going-in position is that all reports just exist in a flat folder or flat directory somewhere and that's why the file name construct is kind of important in what it looks like. Because otherwise, you can imagine making elements of the file name be part of the folder name instead and bury things that way. So, this just becomes an implementation detail about what we think works best.

So, the critical thing here is think about what do you use in your environment? What do you like? What do you expect? What are registries producing? What do registrars want? What's easiest for you? That needs to be part of the discussion. And I'm sure that, again, one of the key distinctions between all of this and the stuff that Jim Gould had out there from before is that there's a lot of flexibility in what he's got. All of this then becomes very flexible. We don't have to answer it here.

It's now the top of the hour. I think it's 15 minutes, right? We'll come back in 15 more minutes, 15 minutes after the hour. So, we'll put a pause in here for those who are remote. Just come back in 15 minutes and we'll see you then. Thanks, everyone.

We are a couple of minutes after quarter after, as these things always go. We know how the breaks are always a little longer. I just wanted people to know we're paying attention. Maybe just a couple of minutes and we'll get going again here.

Okay, if we can, let's get started. So, I'm James Galvin, again, from Afilias. I'm just chairing this afternoon meeting here of the joint REGEXT TechOps meeting, continuing after our break.

So, we ended. We had spent a good portion of our last – at the first half of this meeting going through a registry-registrar report proposal for how to create reports. We left the break with two questions – at least two questions. Certainly, the discussion doesn't have to end here just at the moment, with respect to one question was about the relationship with the dataset file format document and the other question. I can't believe it's just escaping me. One of you guys over here, what was the question we ended with?

UNIDENTIFIED MALE: Is there some middle ground between RFC and BestPractice.domains?

JAMES GALVIN: Right. Is there, basically, a different way? Because we basically have two paths here. We do have the BestPractice.domains website which has a number of specifications on it. This proposal, which is trying to take advantage of the IETF process in particular in order to ... It does two things.

One, it sort of collapses the nine different documents into something a little bit different to fit into the IETF process but it's basically leveraging the IETF as a mechanism for specification.

I'm open for proposals and discussion. I don't know that we can necessarily answer that question right here and now unless you have something that you want to put on the table. I'm certainly willing to discuss it.

I'll just say that my motivation for this particular proposal stands – is derived from discussions that we've had in past TechOps meetings, the last ICANN meeting, and at the Summit where, from my point of view, as a registry service provider – so I'll put on an Afiliast hat for a moment, if you will. That is that I really do require that there be a persistent archive specification. As much as I see the BestPractice.domains website as a good-faith attempt at all of that and it's there, it's risk in terms of being persistent and archival is vastly different than the IETF.

So, whether we like this particular registrar reporting proposal or you like the idea of nine separate documents coming into the IETF, is there something else in the middle? I'm willing to talk about all that but that is the one key characteristic which is essential to me as a service provider, before I'm willing to implement these things. I really do want to implement them. I want to do them. But I can't do it unilaterally. I need to know that there's community consensus on this is the way to do it and then I'm willing to go do that.

There still has the binding question that Thomas has been raising. I don't have an answer for that here at the moment. This is not really a policy group, per se, but if we get that characteristic – that standardized specification here – I that myself, at least others, have informally indicated to me, and I'll let them speak for themselves – I don't want to

speak on behalf of anyone – but they’ll move in this direction. I think that’s what we all want. Roger, you had your hand up.

ROGER CARNEY:

Thanks, Jim. I’ll throw out a proposal. Let’s remove – I don’t know how many drafts there are. But let’s remove the current reporting drafts off the RFC or off the IETF REGEXT. Let’s move forward with Afiliias getting us a draft of one or two RFCs for these two registries, say, by early December. Then let’s start having interim meetings in December and January to work through that set of documents to get into Barry’s hands early next year. Then we can move forward with that. That’s what I would propose. Thanks.

JAMES GALVIN:

Thank you for that. Anthony, in the break, one of the interesting things of substance that had been raised to me about the substance of these documents, I wanted to give you an opportunity to say it in this room, so this group is aware of it. Again, we don’t necessarily have to make a decision about it but it is actually an important detail in terms of what goes in this document and what doesn’t. I liked the question but I’ll let him speak to it.

ANTHONY EDEN:

Sure. What I said was it would be really good to table all of the items that talk about how to deliver these reporting mechanisms. Anything that talks about protocols, anything that talks about naming of the files, the concept of files and directories and all that – and get rid of all that

and focus specifically on the content that goes inside there and the naming and definition of what each of those names and things is. I think if we had just that as a starting point, that would already get – that’s miles from where we are now.

JAMES GALVIN: Okay. That sparks some conversation. Jody?

JODY KOLKER: Yeah. We basically had the same conversation over here. Yeah, we agree.

JAMES GALVIN: Excellent. Ben?

BEN: Yeah. Just as a minor process, suggestion, just as a way to get the first final rough draft that would go into the much more heavyweight processes is just taking these documents and putting them on something like Google Docs or something any other company does that I can just highlight text and put in a quick comment in a few seconds and other people can have a discussion.

Some of the things that need to be resolved are very simple, the order of the date, whatever thing, and the file name or what exact format [inaudible] or things that aren’t important in the grand scheme of things and just need a little bit of finessing and discussion to resolve

them. And just in an online collaborative document editor is probably a better, faster way of coming to consensus from that kind of stuff and these whole drafts of RFCs going through this really vigorous process. That's how we do work. That's how we do this stuff [at my work] at least. Thoughts?

JAMES GALVIN:

I know that ICANN currently finds a great deal of value in [inaudible] editing on stuff and they actually do use the Google Docs and Google suite for doing that.

At the moment ... We'll see. I'm not sure. Part of the problem is you have to put this into what the RFC editor expects and you've got to do this in XML. I think that maybe even the point of view of substance, we'll see where we are and go from there. But I take your point. I use it a lot. Just not sure it's the right thing to do for an IETF document. Roger?

ROGER CARNEY:

Again, I think that [inaudible] take this up and he'll have this done in a few weeks anyway for us. I don't think it will be that much work for this group. Afilius has a little bit of work but I notice that Gould has his hand up. Thanks.

JAMES GALVIN:

Thank you for that. Jim, go ahead, please.

JIM GOULD: You can hear me?

JAMES GALVIN: Yes, we can.

JIM GOULD: It's Jim Gould from Verisign. I sort of wanted to discuss a little bit about the data set file format. It follows on to the comment related to focusing on definition, because one of the big elements of the data set file format is to include all meta data inside the header or in the definition itself. In essence, don't rely on a file name to contain any meta data about the reports themselves.

So, in essence, if you look at the data set file format, it pretty much mirrors what was done for the CSV format for the data escrow, where in essence, [inaudible] key advantages [inaudible] the fact that all of the base objects of the registry that's used within [EDP] is already defined. So, therefore, all the fields for domain host contact, registrar, are already defined within the draft. So, therefore, you really don't need to have an IANA registry for every field. In essence, there is full namespace support and a client could automatically consume them and be able to validate the format using the EPP format that's already defined.

In essence, I just wanted to bring up the fact that if you want to focus on definition, something like the data set file format, that would be updated to be able to support report definition using the fields that are already in there would be able to provide that kind of functionality. Thank you.

JAMES GALVIN:

Thank you for that. Any comments or additions from anyone? I'll respond by saying that I do think it's important for the working group to consider really carefully the distinction between these two things and the benefits of each and decide which path we want to go down.

Are they complementary or do we only need one of them? I don't think that they're straight-up competitive. They're not purely competitive because they each cover different things, although there's a fair amount of overlap there. And I do think that, wearing an IETF Working Group chair hat, it is important for the working group as a whole to think about the distinction between these two things and think about your own environments and what's going to work for you and be prepared to talk about that.

We don't have to resolve that distinction now. In a moment, I'll come back and I'll propose the next action for this as we move onto the other things. But you had your hand up. Go ahead, Roger.

ROGER CARNEY:

Thanks, Jim. I don't think these compete. I think Jim is right. I think the data set file format could be tweaked to make it work. But I don't think we're looking for that heavy of a hand at this solution. We're looking for something more straightforward and simple. But I think the data set file format still has a purpose, just not for these reports.

JAMES GALVIN:

Okay. I think that's good to know and important stuff to ... When we get into a more detailed discussion in the IETF context we're going to ... Because the data set file format document, it's not formally adopted by the working group yet but it is visible and on the working group's radar. So, it's kind of in the same state that this proposal would be.

And it would be fair for the working group to adopt both documents and then figure out what they want to do about them. That's ordinary IETF process. There's no issue there. In fact, I would expect that ...

So, let me bring to a close here the discussion on this document to propose the following specific action at a minimum. There is an IETF meeting in two weeks. I do think that it will be I'm to give a similar presentation in the IETF context, just as we did here, and see what discussion comes from that at that time. I would certainly welcome for everyone here – again, the important thing here is join the IETF REGEXT mailing list. The pointer to doing that was on the very first slide that was up here in this meeting in this slide deck, and if you participate in the meeting, that's great. In any case, any working group adoption decisions are made on the mailing list and be able to pick those up and deal with that.

The action for this is, a bit ago I was thinking that we would produce a document and publish it. I don't want to commit to that just yet. I have to get back and coordinate with Joseph, my co-author in all of this. But certainly we'll have a similar slide deck for presentation at the IETF

Then, Roger just suggested, in a short period of time, like maybe getting into December – so a couple of weeks after the IETF meeting – once

we've had two discussions about it to get a lot of feedback, then we'll map all this out and produce a document proposal. Then, on the mailing list we'll need to have the discussion about data set, file format, and this proposal and we'll figure out what we're going to do going forward.

Then it's however long it takes the working group to resolve that discussion, come to a consensus, and then we can push it all forward and be done with it. Roger?

ROGER CARNEY:

Thanks, Jim. I'll take that action and I'll actually post to the list what I proposed earlier and bring up the data set file format as well.

JAMES GALVIN:

Sounds good. Appreciate that. Anthony, too, I hope that you'll ... Well, I think at the moment ... I think that we'll probably eliminate the file name discussion and the large file discussion and all of that from our draft, I think, based on our discussion in the hallway. You brought up the issue here. So, I think that's good. It moves that line about distribution mechanism versus standardizing all the elements and I like that. So, good thing we'll do that anyway.

Okay. Just as a reminder, our goal here is to consider all of these things and then figure out what we want to try, what we would like coming out of this to move towards adoption in the working group, so that in the IETF working group, to what we want to propose there and do.

We've already eliminated registry mapping. On the one hand, strictly speaking, there's really only four things up there, although the registry-registrar reporting has two documents that would come in but we're only talking about four milestones that we want to latch on to. I think that we could probably convince our area director to let us have all four of them if we would like, if we want to work on all four. But let's finish the discussion of the other three and decide.

So, it may not be that there's any real vote to be had here since we're down to four of them. We can seek to get the document authors to propose to the working group to adopt them and then we'll go from there in setting up milestones and stuff for them.

The next one is the secure AuthInfo transfer. Rick, why don't we bring up those slides, let you walk through that and I'll turn this over to you now.

RICK WILHELM:

Very good. Thanks, Jim. We're going to go through the documents related to the secure AuthInfo transfer and this is an ID that's been out there for a little bit. You can flip to the next slide.

We talked about this this morning. We've got pretty good overlap from the folks that were in the room from CPH TechOps. There are some new folks from the IETF meeting that were in here. But there's an ID and someone will go about pasting the link into there. You can search that up using the terms that Jim Gould and I wrote up to work on a set of best practices proposed to improve the handling of AuthInfo codes

related to – that are used in domain name transfers in the space and how they are handled by registrars and registries.

One of the things that I emphasized during the morning session on this is that one of the things we don't get into is things related to domain name transfer policy or things that happen outside of the handling of the AuthInfos related to the particular form of authority or other things related to the transfer policy. So, those are out of scope. This is just a set of best practices and the ideas written in that form that we're proposing here. So, we can flip to the next slide. And we'll go through these and then have some discussion. Jim, just I make sure that we ... During the morning session, we spent a fair bit of time on this. I want to make sure not to go too long here. What sort of timeframe are you trying to aim for here?

JAMES GALVIN:

I think this is the only other thing of substance. I don't think there's a lot to say about the other two documents. As long as you agree with that, because you know what they are, then I think we've got plenty of time. We're good until 4:45, which is a little over an hour from now.

RICK WILHELM:

Okay. Well, I don't think we need to spend an hour on this. Jody is giving me that look. Thank you, Jody, for that gentle tug on the reins. So, we'll just go through these rather crisply here. And for those of you that weren't in the session this morning please do come and ask some questions if you'd like.

So, a couple of things here on this first page. You can flip to the next slide, please. First off, the AuthInfo, we document in the BCP that it should be produced using defensible, strong – it should be strong and it should be random. We propose 128 bits of entropy using the standard printable ASCII characters except for space to get to 128 bits of entropy. Math tells us that you need to use at least 20 characters. You can't use the space because, in XML, if that space would end up at leading or trailing spot in the position of the AuthInfo, it would be truncated, leading to mass confusion and incorrect matching on the AuthInfo. So, space is not a valid character in an AuthInfo. This is already part of the standard. But when you do the entropy math, got to going to up to 20 characters to get to 128 bits of entropy.

Obviously, those generating the AuthInfo codes would be able to go more than 20 characters but this would be a recommended minimum of 20 characters on the AuthInfo to get to 128 bits of entropy.

Secondly, we recommend in the Internet draft that those AuthInfos be short lived. What that means is that for a particular domain, the normal state of the AuthInfo value for that domain is that it is null. That is, other than when the domain is going through a transfer process, the AuthInfo would be null or empty and then the currently sponsoring registrar would only set the AuthInfo to some valid non-null value when there was a transfer that was in process. Then the client side, the registrar – the sponsoring registrar – would be able to unset it based on its local policy. In other words, reset it to null if it thought that, based on its judgment, that the TTL – it's local TTL – had expired and then it would be able to reset it to null if that transfer window had closed out in its

judgment. The registry would not be – BCP says the registry does not interfere in the set AuthInfo code by a sponsoring registrar, unless of course the domain actually transfers.

Third one here. The AuthInfo should be stored securely. In some cases, what that means is it's not stored at all. The best way to store something securely is to not store it. There's something metaphysical in there, I think. So, the registrar would be recommended to not store it at all. Then the server, if you store it, it would be stored using a cryptographic hash. During the morning, we got a comment – I think it was from Ben – who recommended that we have some language in there about the exact mechanisms about that hash. In other words, that it not be a silly hash – and that's my sort of Monty Python grid, non-technical description. But a cryptographically interesting hash that made it be more secure when it's stored on the registry.

So, those are the first three here. I'll linger here on this slide and take any questions on those. Seeing none, Zoey, we can go to the next one. Jim Gould has put something into the chat that says that it does provide a minimum requirement for has and we'll revisit that to make sure that has some standards there.

The next slide shows the transfer flow when these AuthInfo things are being followed. So, the first two steps are when the domain is initially being created. So, step one is the registrant would go to the registrar to say, "Please, create my domain." Then, the registrar would send the domain create command with the empty or null AuthInfo to the registry. The domain create would follow the sunny day path here in

this slide. So, the domain create goes and the domain is created and then some period of time passes. Whatever requirement that's needed to allow the domain to be transferred would pass between step two and step three and some other stuff would need to happen before the registrant would go and request the authorization info when the registrant would be interested in transferring the name. Registrant requests the authorization info, and then from the losing registrar – so, losing registrar to the prospectively expected losing registrar would generate a secure AuthInfo value using those randomization mechanisms that we talked about, provision it in the registry where the registry would store it using the hashing that we described on the previous slide. Then the registrar would return it to the registrant and not store it locally.

Then, the registrant would be able to then take that AuthInfo code and take it over to their gaining registrar and the gaining registrar would be able to verify that using the AuthInfo, checking that with the info command.

There are a question around whether any registrar could determine if the AuthInfo is set and the answer to that is no. We got that clarified. There's also a question about whether or not if there is an AuthInfo code what the response code would be if you have the AuthInfo code. I think that I wrote down here that that's 2202, that if you're sending it in and you don't have the AuthInfo code, if you send in the wrong one, you would likely get back a 2202 from that.

So, then the gaining registrar would send that AuthInfo code on the transfer request and then assuming that the transfer goes through okay, the registry nulls out the AuthInfo code upon the successful transfer.

Let's see. We've got some stuff going into the chat here. Jody has a question or comment in the chat. In the end, though registrants are not going to read the password and type it into the gaining registrar. They're going to cut and paste it into one interface to another. Do you want to come to the mic, Jody? Go ahead.

JODY KOLKER: I was just commenting on Justin and Jim Gould's I think. I should probably put my glasses on.

RICK WILHELM: Sorry. I did not see that chatting go by.

JODY KOLKER: Yeah. All they were asking is if we should get away from confusing characters like zero and capital O, capital I or I and L. Those types of things. All I was saying was that, most of the time, registrants are just going to cut and paste that password. They're not going to have it read off to them or type it in one by one. That's all I was saying.

RICK WILHELM: Yeah. Thank you very much. I would actually concur with that about doing things to make it – omitting those potentially confusable characters actually would cause you to drive longer length into the thing which, of course, makes the AuthInfo code longer, thereby causing even more difficulty for the poor registrant who now has a 32-character and they’re taking on water quick is I think the gesture that Rich was making there. It sort of gets into a situation. And I would agree, Jody, with your comment in the chat that most of the time they’re cutting and pasting this stuff in and around interface.

JODY KOLKER: Never underestimate, though, their ability to actually deny what you think they’re going to do and do something completely different. The nice thing is that the registrar policy can actually fix that. Internal policy could say, well, we’re going to have longer lengths with fewer confusing characters. That’s fine. It doesn’t seem to block that, so that’s good.

RICK WILHELM: That’s a great point. This really is the “recipe” that the registrar uses to generate the AuthInfo codes can be locally customizable, so if a particular registrar wishes to use potentially – wishes to eliminate potentially confusing characters, they have the ability to do that using a local algorithm. Sort of like your password generator that you have on your phone probably does something like that. At least mine does. Good discussion. Good points. So, we can flip to the next section here. Next slide.

This is all put in within the ... It's all set up ... The BCP has written such that the EPP support is already there within the existing RFC, so in other words no RFC changes are required to support this. We had some discussion a little bit during the morning session about just kind of confirming this but it's written such that it doesn't require any changes to EPP. Null passwords are fine. It doesn't require any extensions or something like that. So, all that support is in there. A set of best practices can be adopted within the existing standards which is sort of the gist behind best practices. And we can flip to the last one on the last slide.

Inclusion. As I mentioned, no need for a new extension. The couple biggest points out of this are really driving home the fact that some registrars are, I would expect, already doing this. The notion of making the AuthInfo is very temporary in nature and only provisioned into the registry for the duration of the window that the registrar seeks to have the name available to be transferred according to its policy. And then when that window closes, which is either per that customer or per that domain, then the registrar would have the ability to de-provision it, null out the AuthInfo. That's what we mean by client-managed TTL. Then, define some best practices for not storing these things on the registrar side and then storing them securely on the registry side using the hashing mechanisms and such. So, fair bit of it is good hygiene but we found it to be an interesting process to go through the process of documenting it.

Justin has a question in the chat. Let me read it, but then Justin, if I don't get it right, you can certainly come to the mic. "Quick question on

the current EPP standards. If the AuthInfo is empty at the registry, sending a null value for the transfer request command will always fail.”

The answer there, I believe, is yes. Jim Gould’s hand shot up and I believe that Jim is here to help answer that. Jim Gould, please, go ahead.

JIM GOULD:

Absolutely. That question came up many, many different times. In essence, if it is stored as a hash in the registry, an empty AuthInfo, for one thing, would need to match the hash. So, that’s not going to match. The other thing is that matching a hash to a null value in the registry wouldn’t match either. So, in essence, passing a null AuthInfo on a transfer request or an info command would both not match an existing set AuthInfo or an unset AuthInfo. Thank you.

RICK WILHELM:

Thanks, Jim. Justin, did that answer your question? “Great, thanks,” says Justin. The stuff that Jim said there is all within the existing transfer mechanisms. There’s nothing really special there around that. But it was interesting how documenting this brought some of that stuff to life. Yes, please, go ahead, Vlad.

VLADIMIR SHADRUNOV:

Thanks. So, quick, question. Firstly, this is wonderful in terms of going forward [inaudible]. But when we look at the legacy inception of this – so, current domain names, for example – so, moving forward as a

registry if we all decide to implement this, that's perfectly fine. Every domain name that is transferred will have the registry nullify the code. We'll generate a new one and so forth. What happens to the current domain names that have AuthCodes at the registry and so forth? And I'm not sure this is the place to have that discussion or not, but essentially, [inaudible] moving forward concept over here coming out of this draft versus a legacy concept where you still have a whole bunch of domain names with AuthInfo codes and registrars still manage them and everything else. Is that to be considered here or do we consider that somewhere else or do we even consider it at all?

RICK WILHELM:

Very good. Good question. So, that is really up to the particular registry, if a registry is going to adopt this BCP what that registry would do. This BCP says that this standing state of AuthInfo codes would be null. So, if a registry was going to adopt this BCP, presumably that registry would make some sort of announcement and they would work with their registrars to identify the transition period and the ramifications related to that. And that would be registry operator dependent and maybe even TLD dependent.

So, this particular draft doesn't define exactly what that transition mechanism would look like but it would be localized in its situation, in its implementation. It would be – if a registry is going to adopt this, part of their messaging, presumably, would be something that would be a positive. We're taking steps to improve our security and stability, so it would be more of a positive message as opposed to something that

would be not positive or neutral. At least that's sort of where we were thinking as we were drafting. Please, go ahead, Vlad.

VLADIMIR SHADRUNOV: Sorry. There's nobody else with questions here because [inaudible] back and forth. So, another question I have, I believe somewhere within here I heard that the AuthInfo will no longer be returned to the registrar as part of the EPP request. Is that correct? Did I hear that right? Because if that's the case, then I'm looking at, from an auditing perspective, as a registrar, maybe we may have overlooked something. Maybe the registry didn't accidentally forget, for example, to null an AuthInfo code after a transfer, something may have happened and I sit here with a domain name that has an AuthInfo code that could potentially be abused by somebody who knows that there's an AuthInfo code on this domain name.

So, removing that element, even if it is just simply to show a null within the PW element, that the main PW within the AuthInfo would potentially raise the question of the registrar, "Why is there something there? Do I need to override it? Do I need to potentially nullify this AuthInfo code? Should I contact the registry because there was a transfer and they didn't do it?" Something of that nature. I'm trying to identify. We're potentially moving a critical piece of information that might lead to certain circumstances just because it has not been blanked or it has not been removed from the response.

RICK WILHELM:

Thank you. It's a good point to highlight. Here's the way I would answer that. It's both ... Since the BCP recommends that the standing state of the AuthInfo code is null, the fact that the AuthInfo code being null is being returned, it sort of effectively means it's right only. So, the registry is no longer returning unencrypted AuthInfo codes because they're not stored – that hash that they're storing it in, it's sort of a one way. It's Hotel California, I hope. Is that my term, Jody? Right? Sure. Jody is no help at all. Not playing the [inaudible] today.

So, yeah. It goes in. It's inbound but it just goes in and doesn't go out. But, that's actually okay because, for the most part, the domain's AuthInfo code is null for most of its lifecycle and it's only a very brief blip that the AuthInfo code is non-null. So, that's kind of okay.

So, the state that ... If there's a registry that's operating where it's been returning AuthInfo codes, that means that it's been operating in a world where they've been non-null. So, post-VCP adoption, standing state of those AuthInfo codes would be null, and so therefore returning that null value is sort of obvious. Does that make sense? Good point.

Okay. I think that's it. Like the note there at the bottom says, please review the draft. Sorry. Justin has a comment. "What's to prevent a registrar from circumventing the TTL? Point one, create a domain with null AuthInfo, immediately set the AuthInfo in passage to the registrar."

So, there is nothing to prevent a registrar from circumventing the TTL because the TTL is registrar specific. It's handled on the basis of by the registrar. If a registrant wants to work with a registrar that operates in that fashion, then they're totally fine. It's a free market and they can

take their business where they want. There may be perfectly valid reasons for a registrar to operate in that fashion.

So, this is, just like there's different security mechanism that make sense for different folks. What we're proposing with this – and that has been, I would say, the situation that Justin has outlined there in the chat with creating new domain with null AuthInfo, provisioning it and then handing it back and leaving it provisioned was sort of maybe the way that this was originally conceived.

But what we're proposing with this BCP is a little bit more modern way of doing it, sort of in the same way that a reset password link that you get from particular website, that link only is valid for 24 hours or ten minutes or two hours or a week or something like that, depending on what kind of website. It sort of has that more of a feel to it. Roger has his hand up and then Owen. Sorry, I'm not sure who was first.

ROGER CARNEY:

Thanks, Rick. Just reading Justin's notes. It's nice because we won't have to worry about policy changes to make this happen. But I would maybe foreshadow that the transfer PDP will probably look at this. I'm sure the question will come up and say, "Well, should the registries have a default of six months, then they should blank it." I'm not suggesting it. I'm just saying I could see that come up later.

RICK WILHELM:

Right. Let me respond to that one real quick and then we'll go to Owen. So, the reason why having a registry-enforced TTL is problematic is that

we will get into these boundary conditions with weird timing scenarios around where the – whatever timeout is picked, if the registry is interfering – and I pick that word very particularly and I keep using it whenever we talk about this – if the registry is interfering in that timeout, where the registry will, at some point, statistically it will happen that the registry will come in and wipe that AuthInfo code out at precisely the wrong moment. It’s just the way that math works. And it will matter and there will be a well-meaning registrant in the middle of a very important therefore and someone’s phone will ring off the hook just because math kicked in. Go ahead, Roger.

ROGER CARNEY:

I don’t think you have to convince us. It’ll be when that argument comes up, remember your response.

RICK WILHELM:

Thank you, Roger. I would also encourage the registrars within earshot to make those points because sometimes if the registry says it – me being Verisign and here being a registry, Jim Galvin being a registry, [inaudible] being a registry – it’s sort of put forth as in a self-serving, “Well, you just don’t want to do the work.” It’s not a technical thing. The problem is what will happen is that, inevitably, it will cause a dispute to arise where no one needs to dispute.

And what I would propose, just so everybody knows, would be something instead of that, if there are TTLs that are out there, if there AuthInfo codes that have age in them, that is over some sort of a

standard, that the registry could notify via reporting mechanism or something else that these are old AuthInfo codes and provide that information appropriately. But for the registry to come in there and wipe them out is sort of a different level of interference. So, thank you, for that. Owen, please go ahead.

OWEN SMIGELSKI:

So, this question is kind of coming a little bit from my background with ICANN compliance and some of the, I don't want to say battles with the registrars, but more discussions, was the concern about how to confirm whether it was the registrant who retrieved the AuthInfo code and when that occurred because that wasn't anything in the transfer policy. With setting the null AuthInfo code and then springing to life when they want to have a request – I'm assuming that could be something that could be logged, which could then in theory be used to further demonstrate that it was the registrant who either activated that or requested that or did something along those lines. That's something that could be contemplated by this.

RICK WILHELM:

Yeah. That's a great point. So, using the provisioning of the AuthInfo code as evidence that a transfer operation was initiated, presumably by the registrant. Yeah. It would trigger residue in the registry that the AuthInfo code was set by the registrar at a particular date timestamp and that at least it would, in our registry, and I would presume in most. It's a great point, Owen. Question down here.

UNIDENTIFIED MALE: One question to the TTL. Can you explain me the difference in the dispute if the registrar nulls it at a specific time compared to a registry nulls it at a specific time? I don't really see the difference. If you make an alt ID like November 1st with a 14-day TTL it will expire, November 15th no matter if the registry nulls it at that time or the registrar. It will be the exact same outcome.

RICK WILHELM: Thank you. Yeah. So, the difference is that the timing of the expiration should be done by the currently sponsoring, AKA the losing – prospectively losing – registrar which has visibility into the appropriate TTL on the AuthInfo code. The registry doesn't have the context or the perspective about when the right time is to expire that.

So, the registry, if it comes and expires the AuthInfo code – so, good question – by wiping it out, it's doing so just based on some sort of a timer and it doesn't have the localized knowledge about what's going on with the particular customer.

So, for example, a registrar might issue an AuthInfo code, tell the registrant that it's good for seven days, and then make the unilateral decision to let it go eight, or based on something else that's going on in the account, might wipe out at four days. But that's a decision that the registrar can make based on what's going on with that particular customer in that particular context. And the registry doesn't have any of that awareness, and if it just so happens that the registry wipes out

the code at the wrong time, the registry is breaking a transfer that was in process and could cause service disruption.

UNIDENTIFIED MALE: I don't see any problems with that registrar could null an AuthID at any point if they see a need for that. For example, if you have an eight day and you feel that you need to null it after four days, that's perfectly fine. But I also don't see a problem with having a standard max time of 14 days, for example. That could be just a general standard registry level. After 14 days, it nulls. Then a registrar could still null it before that.

RICK WILHELM: Yeah. All that is possible. But the 14 days is not ... Whatever number we pick, there's perfectly practical business reasons why a particular registrar might want to allow it to go longer. So, there's not any one number that's more magical. And the registry doesn't have the local knowledge to be able to make that decision and it also causes the registry getting in the way of the transfer between the two registrars. At the boundary condition where it's most critical.

UNIDENTIFIED MALE: Just out of a security standpoint, I don't see any argument valid of why you would have an AuthID still working after like six months. I don't see you can make that argument. You can always make a new AuthID after 14 days, three weeks, whatever. But I don't find any argument on why we should allow an AuthID to be valid after six months, for example. I would love to see the argument that would work in that case.

RICK WILHELM: I think over here first and then Vlad.

JOHN: John from Name.com. I could see the argument being that if you have a customer with a lot of domains that likes to manage all of their domains and know what their auth codes are so if they need to move that they would want their auth codes easily accessible. If they have more than a thousand domains, it ends being a lot of work. So, in that case, I could see that type of customer wanting that and wanting those auth codes available and valid all the time.

RICK WILHELM: Vlad?

VLADIMIR SHADRUNOV: So, question to your point. This, to me, skirts the lane of policy to come into play where you have been given ... The registrar has provided an instruction to the registry to perform a certain task. Then the registry via its policy is overwriting this action, essentially, which is perfectly fine as long as I understand that that falls within the [inaudible] of the registry's policy and its own operations in that sense.

So, changing or undoing an instruction – and I'm not sure on the legalities of this but typically, to me, that falls of one should not change the data that has been provided because that data has been provided with certain recourse and with certain requirements and actions. The

amendment of such data would then fall, as part of the policy requirement, such as every single successful transfer, the registry is mandated to null the auth code, for example.

That's one of the issues that I kind of have with point seven in the process because that, to me, seems very much like a policy statement rather than something that would be functionally because you're including both parties. You're leaving up to the registrar to set the auth code. You're leaving it up the registrar's internal policy for how long that will last. Why are you not leaving it up to the gaining registrar in that sense to also reset it after the transfer? Why are you putting the registry, essentially a policy statement for the registry to perform a function – a technical function but still a function? That will also then extrapolate up to, let's say, six months or three months or whatever it may be. That, to me, is skirting between a technical implementation and the lane of policy. In my opinion, at least.

RICK WILHELM:

Right. A couple of things there. Just to pop the stack here, because we've pivoted into a slightly different topic there – and I'll come to that in a second.

Just to be very clear, this is a BCP and it does not set policy, nor is it a standard. It's a BCP. So, I just want to be very clear. So, we're talking about what should the maximum be. Within the IETF, within an IETF, is not a standard and is also ... So, that would not be within its remit. Whether or not ICANN would choose to set a policy regarding maximum length of time is sort of a different discussion entirely.

And to your point about the puts and takes of should there or should there not, another reason I think your point about certain registrants wanting to have their auth codes provisioned because that's sort of – domain investors are ... I won't call them necessarily investors or whatever – people with large portfolios – that's a very objective term – might have different ways how they want to manage them and if they want to leave them ready provisioned with the AuthInfo codes.

Additionally, different registries may have different maximum TTLs which could have it very difficult for registries and how those get provisioned.

So, within the IETF, it would be staying away from the notion of what that maximum would be. So, just to be very clear about that. Which is why there's nothing that's stated in the BCP.

Regarding the BCP statement of don't keep the same AuthInfo after transfer, the reason that's in the BCP is that the transfer is no longer in process, so since the transfer is no longer in process – and one of the things that the BCP says is that if a transfer is no longer in process, the default state of the AuthInfo code should be null, that's why it's reset. So, that's the thing. It's really “a convenience” for the registrar to wipe it out. So, that's why it goes to zero. Please, Anthony.

ANTHONY EDEN:

Could it be possible just to change that just to indicate that it is wiped out instead of setting responsibility or saying that either the registrar or registry may wipe it out? Because I can see where you're coming from.

It doesn't have to be the registry. It doesn't have to specify. The registrar could absolutely do the gaining on completion as well. I don't know.

I mean, it is in the specification that defines the registry is doing that operation. It may not be necessary. Just a small tweak.

RICK WILHELM:

That's a good point. We'll take a look at the exact language around there and maybe perhaps capture the rationale to help get some of these topics in there. Thank you. Hopefully, that's helpful.

Question. Jim is pointing at Rich.

RICH MERDINGER:

Two things. One was I wouldn't ... This goes back to – oh, my God, I'm tired – the gentleman that is no longer here's comment about ... Owen. Thank you. Inferring things from the request of an AuthInfo implying that the registrant intended anything because an email account can be compromised. You don't know that the registrant initiated anything. You know that the registrar has initiated the request and that's all you know by this.

Now, it can be used for a lot of good reasons, a lot good purposes for looking into the issue. But I just don't want people to think that they can infer who did what based upon the fact that an action was taken by a registrar because someone may have compromised an account.

RICK WILHELM: Okay. I think you had two points.

RICH MERDINGER: I did, and now that Owen is back, it's gone. So, there you go. If it comes back, I'll interrupt you again.

RICK WILHELM: Is that an implicit request to make Owen leave? I'm not sure if that's allowed. He's paying me to do it. Yeah. That's an interesting [inaudible], Owen. I like it.

I think that when the ink would actually hit the page, I think the text would acknowledge the point of what you said about the registrant versus the registrar and who actually indicated what, etc. Saying it's a good refinement that you're highlighting there. That's a very good point.

Okay. Other questions or comments? Vlad, please, go ahead.

VLADIMIR SHADRUNOV: There's a whole bunch of these things coming to mind now. Just to bring in some more policy notions here, in the current IRTP I believe there's a statement that says that, one, after a successful transfer, another transfer cannot be done to another registrar for X amount of days, except to the previous registrar.

So, in this case, when the first transfer is successful, would [inaudible] any registrar have to re-update the auth cod and pass it back again to

the registrant a new auth info code in order for them to give it to the new – now old or losing registrar – to reissue the transfer back to the previous registrar because there was potentially a mistake in the transfer? So, the whole process, would it be more complicated in that sense? Is it even an issue? Because I'm just looking and going, "Okay, so I have a domain name I'm transferring from registrar A to registrar B. I have my AuthInfo. I'm perfectly fine. Oops. I made a mistake." Whatever. Something internal happened, whatever may have happened. We would actually want to go back to the previous registrar. I can just use that same AuthInfo, for example, currently and just transfer it back and it was updated and so forth. In this case, the whole process of re-updating, re-issuing, and redoing the transfer would essentially occur. Would that be the case?

UNIDENTIFIED MALE:

Based on what you said, Vlaid – and I'm a nefarious character – I'm going to sell you a domain name. I'm going to give you the AuthCode. And as soon as I got the money, I'm going to pull it right back to the registrar because I still have the auth code.

I think of best practices, or whatever the BCP is, to have that thing reset and go through the overhead in the not-too-common instance where a reverse of the transfer needs to be done.

RICK WILHELM: I think that's a fair point. If you're in a transfer back, snap back situation, resetting the AuthInfo code is the least of one's worries. Other questions or discussion? Quoc? Sorry. Then, Ben in the corner. Quoc?

QUOC PHAM: If we were to implement this BCP, how much notice do registrars need? Just an open question. And does any or do any registrars envision any disruption with this implementation to current operation?

NEIL MCPHERSON: It definitely comes back to this how we deal with legacy registrations. I think that's ... The new registrations were already pretty easy but how we deal with our millions of legacy domains.

BEN: So, just looking at this last slide, it says there's no EPP extension necessary and the TTL is managed on the client side. So, just to be very explicit, that means if the losing registrar creates the code and they decide, "I want this to be valid for one week," they're the ones responsible for [inaudible] that out a week later?

UNIDENTIFIED MALE: Correct.

BEN: Okay. That's not my preferred ideal thing but I think we're maybe past the point of changing that. But I guess sort of as an overall thought, I

definitely think this is an improvement and certainly more secure than what we currently have. So, generally, I'm overall favorable on it. The previous process is, what, 20 years old or something? I think it dates back to the very beginning of this and was never changed. I think this is definitely an improvement.

RICK WILHELM: Vlad, go ahead.

VLADIMIR SHADRUNOV: So, to Quoc's point, from the registrar perspective, the whole process now changes. Whereas previously I might have my own security process in place that would update the AuthInfo once every day, once every two days to the registry and so forth, no longer – well, I potentially do but it's not going to be setting it correctly. It's not in [inaudible] of the implementation. That's BCP. So, I'm going to have to make changes to my process that now if my registrant wants an AuthInfo, I'm going to have to no longer just simply do a domain info and get it or look at my local stores, database, [inaudible] send through an update, and then from there make sure that I complete my logs [inaudible] information from the EP request going back and forth, and I display the correct [inaudible], implement things like TTLs so forth.

I'm not sure. I mean, this is essentially ... How long is a piece of string in this case? Whether you have a small registrar that has difficulty implementing this versus a larger registrar that may or may not have difficulties implementing such a solution and so forth. So, it's really up

to the first registry that would do this, alongside the first registrar that would do this to kind of set that standard, or at least set that timeline of, okay, let's do this, let's see how you guys [inaudible]. So, you, as the registrar, you start implementing this over here. Let's see how long you guys take. Then we can do what we have to do on our side as the registry. Let's see how long we take. And kind of set that, at least for a first-off timeline of this is how long the first one took. With other implementation and so forth, take it from there and say maybe it's 60 days, maybe it's 30 days, maybe it's 90 days, however long it might take.

But because there's development involved, it's kind of an unknown, essentially, until we actually review what needs to be implemented and what has to change.

UNIDENTIFIED MALE:

Just real quick. As far as I can tell, though, since it's best practices, it's not conflicting with your current processes. There's nothing that would break a current process. You just wouldn't be following best practices, until you chose as a registrar to implement those best practices, if I've understood correctly.

VLADIMIR SHADRUNOV:

So, the problem there lies when a registry implements it, and me as a rather, I don't implement that. So, let's say, for example, Afiliat [inaudible] follows this and they nullify all of the AuthInfo codes in their database. At that point, if I haven't got on board with this best practice, then I couldn't do a domain info and get the AuthInfo back and I can't

use the current one that I potentially stored and hashed or whatever in my local database and give that to my registrant to go transfer the domain name. I have to apply an update with the registry to generate that AuthInfo and then go from there.

So, as soon as one registry does this, everybody else as a registrar essentially has to get on board to make sure this is right and it's a really big, big step forward for anybody that does this first. I'm just putting it out there that this is really something that is quite massive, because as soon as somebody does it, every other registrar has to get on board to do it, I think.

JAMES GALVIN: Rick, can I just ask you to bring this to a close here?

RICK WILHELM: Yeah. We can come back. But go ahead, Anthony, and then I want to just comment on this adoption thing.

UNIDENTIFIED MALE: This is a point for making the registrar responsible for nullification because if you remove the registry's responsibility, then you would have the old auth code and you have been given exactly. So, that's argument to leave in the registrar's responsibility.

RICK WILHELM:

Yeah. Thank you. We'll wrap it up here shortly. As far as adoption and transition, the first step that a registry would have to do is be able to support some key features like allowing the create with an empty auth code. The registry has to be able to support that. Also, wiping out the AuthInfo on the transfer. Those are very registrar-facing features. The hashing thing is an internally thing that the registry either does or doesn't do but it's sort of an internal thing.

Then, once the registry does that, then it can be [proceeded] because once that happens, about allowing the null AuthInfo on create and wiping out the AuthInfo on transfer, then the registrars can be able to being adopting that.

Then, as far as notice periods, it depends on your contracts around the notice periods. Conservative implementation, if you think that this changes the API, that's 90-day notice and stuff. But it's still optional. If a registrar wants to provision an AuthInfo code, they could still keep provisioning an AuthInfo code. It's still not have to do that – sorry.

NEIL MCPHERSON:

My question is how far can a registry be running the current practices and in parallel have these new best practices going on? Can both be done at the same time?

RICK WILHELM:

Yeah, I think so, because the registry could be storing the AuthInfo code in a properly hashed format. It could be allowing registrars to provision the AuthInfo code as null. It could be wiping the ... The wiping of the

code upon transfer is either do or do not. That would affect probably all registrars. But a registry could have some registrars that provision AuthInfo codes on the create and some that don't. And those could coexist. Does that kind of make sense?

NEIL MCPHERSON: I'm just wondering, on the rules for the length of the auth code, for example, that would be difficult to say yes for some domains there is a rule and for others, there isn't.

RICK WILHELM: Yeah. That would be a registry policy where the registry would have to decide how they were going to enforce that. There's a number of ways that could be phased in. It could be, look, starting here, we're going to allow you to do this and then, at a certain point, you're going to be required to have at least this much entropy in your AuthInfo codes. That sort of things.

Is that good, Jim? Good discussion. I think this has been a great discussion here, both this morning and this afternoon. Jim Gould and I both really appreciate it. We're going to be taking a bunch of feedback and incorporating it. I'm certainly available here all week to take hallway discussions. Thank you. Thanks for the time, Jim.

JAMES GALVIN: One last question to you in the context of this being a joint REGEXT-TechOps meeting. Our goal here is to get two documents that are

eligible or interested in working group adoption. So, my specific question back to you is what is the status of this work?

RICK WILHELM: So, it's an ID that's out there. We'd like to propose it be adopted as a milestone.

JAMES GALVIN: Okay. That's good to know. Thank you. That's the point there. So, let's go back to the full agenda.

So, there were two other documents. I'm going to lean without any preparation on the principal authors of these other two documents to say a few words about them. You knew that was coming, didn't you? Or you were hoping it didn't.

In fairness, we have the Registry Maintenance Notifications and Unhandled Namespaces. They're both very stable documents, very table discussion. So, Jim Gould will get the last word on his unhandled namespaces. The three of you are sitting over there on the registry maintenance notifications. Which one of you is going to put your hand up?

UNIDENTIFIED MALE: I'll wing it and you owe me a beer. Registry maintenance, I think, most other registrars are familiar with what it is. Instead of 1000 registrars and three registries, now we have 1000 registrars and 70 registries. The amount of maintenance notifications that come in on a daily basis is

several, at last – probably dozens of, “Hey, there’s going to be maintenance on O, T, E or dev or productions.” Well, somebody has to manually read those, get it out, understand what’s going to be down or up at that time, and then send out something to their operations team to say, “Hey we’re going to be down. Don’t call me in the middle of the night when you can’t get a connection to the registry.”

What we would like to do and what this was meant to do was to make that into an EPP poll event so it could be all automated and standardized so that every registry would do it the same way. Then we can automatically read those pull messages and then set something up in our system so that our operations teams would now that they’re down, don’t call anyone. Registrations are going to back up. Name server changes, etc., are all going to back up. That’s what that draft was for.

I’m not sure. It’s been out there for a while. I think it might already be implemented. Has it been implemented by anybody yet? No? No, we’re looking for somebody. Anybody?

Anyway, that’s the whole purpose of that. When we have the second round coming in 2021, there could be more registries out there – 2023, whatever. More registries and more problems. It just becomes very hard for registrars to keep track of that. So, that’s what it’s about.

JAMES GALVIN:

Go ahead, Roger. I’ll let you jump in.

ROGER CARNEY:

I just wanted to add that this isn't trying to supplement anything, reporting requirements or anything registries have already agreed to. They send out maintenance notices and everything via email. This is just allowing us to consume it electronically so that we can process it in a more efficient manner.

We're not trying to change your reporting requirements or anything. We just want an electronic way of doing this. So, thanks.

JAMES GALVIN:

What I want to add to that, highlight what Roger is saying from a registry perspective in a different way, we are registries obligated to provide notifications about a certain set of things. The intent here is not to change any of that. You said it's not a supplement but I want to phrase it as it is actually a supplement to what we're already doing which is it's providing an automated way to manage this so it can be consumed in an automatic way by registrars.

There was actually a separate step here on the ICANN side to deal with if we want this to be the only way that that happens versus other things. That's a separate action. This is just a technical specification for how to provide that data, that ideally registries will adopt and then registrars can start to consume and then we can figure out the policy needs after the fact that goes with that. Rick, go ahead.

RICK WILHELM:

Has there been any – and I'm sorry because I've had my head in other things. Has there been any discussion around this being just as good as

the other contractual notification that [inaudible]? Because we put a lot of work as a registry operator into these registry notices on the maintenance notifications and things like that and they're done to fulfill particular contractual requirements. A lot of attention goes into them for that reason.

So, if there as a way that they could be a substitute for that, such as we did this thing, then it would not require the need to do the other kind of notifications, I think then that would be a game changer. I think that would be – to make it be not a supplement but instead a substitute, as Roger [inaudible] said.

JAMES GALVIN:

So, I'll respond from my point of view, speaking from an Afilius position. What's interesting is I would love for this to be automated. I think this is a great thing. We should move in this direction, and no there's been no discussion about replacing those things.

But I make the following observation that worries me a little bit. There are registrars who don't consume pull messages. That becomes the fundamental problem. There are issues on the registrar side about dealing that. Even if the registries want to go in that direction, that's the other half of the policy side that has to be dealt with. It is what it is.

But, I like it, Rick. From my point of view, I'd love to move to this altogether and automate this stuff. I know there are registrars who would like it. So, for the moment, I see it as a supplemental thing and I

don't know what to do on the back half about registrars who don't consume them. Go ahead.

RICH MERDINGER:

Quick question. Does the concept you people are discussing here include the idea of modification of notifications and revocation of them? So, I guess a modification could be a cancellation. But people read emails now and, "Oh, wow, that changed from 2:00 to 4:00." Another pull message comes out, we might think we have two maintenances kind of thing.

UNIDENTIFIED MALE:

good point, Rich. Yeah. Just to answer Jim's question on some registrars don't pull their pull messages. They're probably not paying attention to their maintenance noticed, either.

JAMES GALVIN:

Yeah. I wasn't really going to go there, but you're right, there's a lot of ways to look at this problem space in which I think the solution is. We can certainly decide that this is the right answer and everyone else is just going to have to come along or their situation doesn't change kind of thing. And that's fine. We can ... We'll deal with all those issues later.

From a technical point of view, this is just about whether or not this is a technical specification to move forward. I'm gathering that you would propose that this be on the list of things to adopt and we move that

forward and then we'll sort of have those larger conversations about what problem is being solved or not in the IETF context. Go ahead, Jody.

JODY KOLKER: I'm just going to answer Rich's question. Yeah. There is a maintenance update in there where it would update it so we would know.

JAMES GALVIN: Okay. With that, let's jump to the unhandled namespaces, also relatively straightforward and stable document in its space. Jim, I haven't actually looked to see if you're still with us, but if you are, would you please just start talking. I'm going to reach out to you to speak about your document there.

JIM GOULD: Yeah. I'm still here. This is Jim Gould from Verisign. Actually, I'll provide a little bit of a background on this particular draft was that this came up during the work on the change pull draft that became an RFC. There was a question raised whether or not the change pull should be inserted in the pull queue if the client doesn't support it.

In essence, the concern there was that if the change pull was put in there, then the client came in later on, logged in and then the login services did not include the change pull message, what do you do about that message?

So, as we were working through that problem, actually Martin had the idea here that – I worked with him on this particular graph to come up

with a solution that was RFC compliant. In essence, the concept was to allow for the information to be retuned while not breaking the RFC and not breaking any [inaudible] that were to validate responses coming back from the registry.

In essence, this particular draft would be able to support pull messages as well as regular responses messages that it could include extensions, that the client did not explicitly specify that they support [inter] login services.

So, this solves a gap that we currently have and that's why I support for it to be adopted by the working group.

JAMES GALVIN: Thanks, Jim. Any comments or questions from anyone in the room? Rick, go ahead, please.

RICK WILHELM: I will offer that this is a remarkably esoteric topic and it takes more than five minutes to get one's head into it. Compared to the length of time that we spent going over a relatively simple document like this, this is one where these unhandled namespaces to really appreciate the motivation for why this is important. It takes a fair bit of presentation to get into it. I support this also being adopted, because during the work with Jim and some other things, I've come to appreciate why this is important, because without it, your code sort of gets to a spot and then will need to kind of throw up its hands and you end up with some sort of a problem that your code doesn't know how to deal with. This gives

the code a landing spot because it defines the spot of what happens when there's no clear path to handle content that you're getting over the pull messages. Thank you.

JAMES GALVIN:

Thanks for that, Rick. It really is a very deep issue. My recollection is, although I don't remember when, there was quite some discussion in an IETF meeting at one point in the past about this, digging into it and some discussion on the mailing list, too, quite some time ago when this stuff first came up.

And you're right. It's a very detailed kind of deep, technical issue. If you're not deep in this, you would probably totally overlook this. But it is an important thing and it's useful to have this. Roger, go ahead.

ROGER CARNEY:

I agree. I think this should be ... I support this coming into the group. It is something – and again, as Gould mentioned, this just came up because of some other we were working on and someone realized there is a gap and trying to describe that gap was kind of difficult. It did take a lot to do. So, I do think it's important. So, I do support it.

JAMES GALVIN:

Thank you. Any other comments or questions? Okay, then let's jump to the next slide here, which is just an indication that we're going to talk about the priority and let's jump to the next slide here. Oh, look, it's the

same one all over again. Stay there. That's intentional. Just a reminder, just trying to keep the slide deck in line with discussion here. That's all.

So, we started ... The objective of this meeting, again, was about trying to walk through the backlog, if you will, the technical backlog that we have of issues and give due consideration to what really ought to be a technical standard in the IETF as a proposed process for making this a technical standard. The vehicle by which we would accomplish that particular goal.

We started out with five proposals. We sort of talked about all of them. We discharged the last one, so it's not under consideration. So, we're left with four – registry, registrar mapping which does not actually have ... It has one document. There's another document to be drafted. The action that I took away here is that we will have some discussion at the next IETF meeting about this, too, to get some exposure there and see what discussion we can have. And we'll produce a document after that which we can then seek to have adopted by the working group. Then we'll have a detailed discussion about dataset, file format, and the proposal that's presented here and we'll sort out what's going to happen there and what to publish.

There's the secure AuthInfo transfer. I think there's a lot of support for that. We clearly had some discussion about some details and that's fine. But I think it's appropriate to have it adopted and proceed and the discussions can continue about the details and we'll sort out what those are and how soon we want them.

The registry maintenance notifications and unhandled namespaces. We didn't have a lot of discussion about them, but again, I think they're both fairly stable. So they're actually fairly ready to move forward fairly quickly. Unless something gets identified and there's a showstopper that comes up, which is always possible. You never know when you get more eyes on things. But we can adopt those and seek to adopt those and see if we can push those along. That leaves us with four.

Technically, well, as a matter of principle – it's not really a hard and fast rule – we only have three milestones available to us to adopt in the working group. But what I would suggest is we can actually adopt all of these – we can seek to adopt all the documents and then we will need to have a discussion about what milestones to put out there and what the timeframes should be for them. So, we may not set four milestones but we'll see what the discussion gets for us in the working group. We'll seek to adopt them. After the next IETF meeting – we'll put all of these on the agenda for the next IETF meeting, so we get some discussion there. We'll seek to adopt them formally in that process and then we'll have a discussion about milestones for them and move them along. And we'll pick and see where we end up with those, how quickly we can do that.

So, while I had thought that we would have a discussion here about priority order, I'm open for suggestions. I think my specific proposal for this group is I think at the moment it's okay if we consider these all equally and let's just seek to move them along and we'll see what happens in the discussion in the IETF. Is there any reason not to consider them equally and just move all of them forward for at least the

document adoption? Anyone have any questions or comments about that? I'm not seeing any hands in the room. I'm getting nods around here. Oh, you've got a hand up, Jody, or just voting in favor? Okay, that's fine. Hand. Go ahead, please, Anthony.

ANTHONY EDEN:

Just, secure AuthInfo transfer registry maintenance notifications and unhandled namespaces are all either done or very close to done. To me, those are no-brainers. I think we're still missing something on the first one which is the final, a single one or two proposals for creating the registry. So, that's the only one that I'm in question about, if we can get it in time for the IETF for adoption, because there's not yet anything to adopt. I feel like that's where we sit.

JAMES GALVIN:

So, the specific step of adoption and call for adoption won't happen until after the IETF meeting. A specific call. So, I think that we can have a consensus here that our objective is to seek that. The actual action and step of doing that will come after because we have to give these things exposure in the IETF working group for the folks who are only there – because there will be some of them who didn't come to this and only there. And ultimately the request actually goes to the mailing list so even more people can be part of it and make sure that we get coverage from everybody which is another reason why I want to remind everyone please join the mailing list. And I see my area director has got his hand up and I definitely want to give him a chance to speak about how this process is going to go. Thanks.

UNIDENTIFIED MALE: Actually, I don't think I'm speaking as an area director here, but just as somebody who frequents working groups. Why don't you put out the call for adoption by email now and have it close after the IETF meeting, so that you don't have the extra delay?

JAMES GALVIN: Excellent suggestion. We can certainly do that on the three because there are three solid documents that are there for those. Well, unless ... I think, Rick, you were talking about there might be some things you want to do to the AuthInfo transfer, right?

RICK WILHELM: It would be just tune-up. I don't think it's any material changes. It would be just fit and finish, polishing, standard feedback that we get. So, I wouldn't consider that as needing to delay it for this. Thank you.

JAMES GALVIN: And I'm going to look over here to the three. And to Jim Gould, I'll look to you online in a moment, too. But the family of three over here, is there any reason not to call for adoption [inaudible] document?

UNIDENTIFIED MALE: No. We've been discussing. Why don't we get the three that are almost, that actually have a document, get those in, get them in there, start

working on them, while we're waiting for the fourth one? The reporting one.

JAMES GALVIN: Okay, thanks. And I see Jim Gould in the chatroom saying that his document I good to go, too. Roger, go ahead.

ROGER CARNEY: And just to go along with what Barry is saying, I'm going to publish to the list our discussion that we had today on the reporting stuff to make sure that that gets solved and there's a path forward, at least, that we can talk about. I don't know if we'll get to an adoption before or whatever, but at least get it on the list that it's going.

JAMES GALVIN: Right. There's no document to adopt for the new stuff that was presented today and I would say that my plan is to open it for discussion at the IETF meeting and then we'll produce a document after the fact and then seek to do a call for adoption. The timing of this will be interesting but it should all work out. It's all fine. Vlad?

VLADIMIR SHADRUNOV: Quick question. What do we mean by priority? Because I'm kind of getting confused based on what I'm hearing. So, we can do all of these at once, right? And if we choose to prioritize, what happens then? Let's say we go with the secure AuthInfo transfer? What happens to the other ones that are below that priority? Do we just do one at a time, wait for

that to finish through the IETF and then go? If that's the case, then just go do them all. If you can do them all, do them all. That to me just makes sense. Just get it [inaudible].

JAMES GALVIN:

So, in principle, yes. These things can all happen in parallel because they're happening on the mailing list, and as long as you maintain email threads and you keep them separate, each can have its own thread. You can move them along all together. Then the administrative actions that go along with it is just up to the chairs to keep up with that and make that happen and push things along.

As individuals who are participating, if you have several documents that are of interest to you but you can only work on one at a time, then that's the way that will drive progress on the things. If you only contribute to one thread and the other one gets held up because you're not contributing, then that's what happens to the documents.

What I think is ... As part of the IETF discussion, we will have to have a bit of discussion about what we want for the milestone, and what I mean by milestone is what the working group will have to decide when it adopts the document, it will have to decide if it wants to ... I can see the last three. Maybe secure info transfer. Maybe not right away. But registry maintenance notification and unhandled name spaces strike me as two documents where we can have a fairly near-term milestone and the milestone is submit for publication where it then transitions to the [ISG queue work queue] because I don't imagine there's a lot of discussion there. They're fairly stable documents.

We might give secure AuthInfo transfer another month in the period just because there will be a little bit of discussion. The reporting stuff probably a little bit longer because there's going to be I think a fair amount or more discussion about that, but we'll see.

So, we will have to pick the milestone and the milestone to pick is when do we want to commit to submitting it to the ISG for publication? That's what that means in an IETF context.

Then, as Barry had said earlier, that starts and triggers a three to four-month process when you add all the steps that have to happen in all of that. So the working group has its own delay for what it wants to do there.

We're actually two minutes by my count past our working group time. I will allow ... If you go one more slide, I think, after this, that takes us right to any other business is the only thing that's left. I think we have a clear set of actions at this point for what we did.

From my point of view, this has been an enormously successful meeting and I hope that people see it that way. That was the intent here. Productive. Trying to move along the work that we're doing in a way that's useful to all of us, the industry more generally. At least I feel really good about what we've done here. I hope that others do, too. And you can beat me up out in the hallway if you disagree with me, so that we never do this again.

But looking at Mark and Tobias as the chairs of TechOps, this kind of relationship I think is a good one and we should take advantage of this.

So, look for opportunities. I'll give you, Mark, an opportunity to speak to that. Please, go ahead.

MARK: Thanks, Jim. I just want to jump in and thank you for volunteering to do this. I think your approach to Tobias and I at the GDD Summit and suggested this and I think there was a lot of support for it and I know there's a lot of work that went on on your part to make this happen. So, I just wanted to take a moment to thank you for taking the initiative, the suggestion and making all this happen today. So, thank you.

JAMES GALVIN: Thank you for that, Mark. I appreciate it. We'll do this again. We get to figure in when we need these kinds of things. So, let's make it happen again. So, with that, I will say that we're adjourned. Thank you to Zoey.

UNIDENTIFIED FEMALE: Thanks. You can stop the recording.

[END OF TRANSCRIPTION]