ICANN86 Seville | PF – SSAC Work Session (3 of 3)
Wednesday, June 10, 2026 – 11:45 to 13:15 CEST

KATHY SCHNITT

Hello and welcome to the 3rd of 3 SSAC Work Sessions. My name is Kathy and I'm the participation manager for this session. Please be advised that this session is being recorded and is governed by the ICANN Community Participant Code of Conduct, the ICANN Expected Standards of Behavior, and the ICANN Community Anti-Harassment Policy.

Regarding participation for today, this session is designed for the internal work and discussion of SSAC members. Observers are welcome to watch the work party conduct its work, but this session is not open for observer participation. For my SSAC members, all members will be promoted to panelist status in the Zoom room. To join the speaking queue, you raise your hand in Zoom, please, and if you're physically in the room, when called upon, state your name for the record.

You will use the Zoom chat with one another, but please note this is visible to observers in real time. For observers, the speaking queue is limited to SSAC members and the observer chat has been disabled for this session. I will now hand the floor over to Chaoyi Lu.

CHAOYI LU

Okay, good morning, everyone. I suppose we start from the DNSSEC operational considerations work party. We have half an hour and we have several items on the agenda. I believe we have some slides. Do we have them? Yeah. So, am I expected to bring that up from -- okay, please do that. Thank you. Okay, next slide, please. Kathy, can we have the next slide, please?

Okay, so some items on the planned agenda. We'll first have an overview of the progress that we have. We have put together some early findings expected to be in our report and we could have a discussion on that. And then we will review what is still remaining before publishing our report and then we'll review the timeline again.

So, next slide, please. So, here we have the overview of progress. We have seven sections in our report and we have made some progress lately and we are still on the way in finalizing especially sections four and five.

Okay, next slide, please. So, here I've listed the outline of report which lists the sections and subsections that we now have. We start from an introduction followed by some educational material about the functional elements including the benefits and utility of DNSSEC functional elements with signing and validation. And then we have this current state analysis section presenting some data about DNSSEC deployment status, about analysis and a summary of outage events and their causes. And then we do a review of

current status of DNSSEC software tooling and automation and how they help in a seamless deployment.

And then in sections four and five, we summarize some determinants for DNSSEC signing and validation and identify the potential issues we're having from these aspects. We discuss the different deployment scenarios and risk profiles in which scenarios will DNSSEC be required or considered more useful and we discuss technical determinants and non-technical determinants and finally some security determinants. And finally, we'll come to findings and recommendations. So, this is the outline of our report.

Next slide, please. Okay, so this one lists some of the early findings we have from our investigation and from the text of a report. So, first we want to say that registrars play an important role in promoting DNSSEC deployment. For example, about their availability of DS uploading channels, supported algorithms, automation and tooling and we call for more efforts in these specific items from registrars. And the second is tooling and automation. Essentially, we find that tooling and automation have significantly offloaded DNSSEC operational complexities, and so operators are encouraged to make use of them for easier deployment and to avoid misconfigurations and outages.

And third one is about financial incentives and other non-technical issues. So, financial incentives on all non-technical issues are notably effective and we also discuss other non-technical issues as we find that DNSSEC deployment is more than just a technical

problem. And finally, we find that impacts vary by context. We find that domains with many delegations or high-profile domains can have a greater impact in this issue.

So, here is a very brief summary of our early findings. We have a detailed version of findings and we can come to that if we still have time. So, do we have any comments or missing items that we think should be findings in this report? Yeah, Peter, thank you.

PETER THOMASSEN: Hello, Peter Thomassen. So, I've been maintaining a small list of registrars with whom our users at SSAC have made the experience that they do not support provisioning DS records even for gTLDs where they are obliged to do it. So, maybe we can call that out. I mean, not necessarily particular names, but maybe we can call out the fact that such registrars exist and then maybe also have a recommendation to, I don't know, lift that up to some higher enforcement level, not in particular cases, but to be more vigilant around this topic.

CHAOYI LU: Okay, thank you, Peter. Are you suggesting that some help is needed from registrars?

PETER THOMASSEN: I don't necessarily know if that is help that is needed. The registrar accreditation agreement says that registrars for gTLDs are obliged to relay DS record updates that are given to them to the

corresponding registry, and some don't do that. So, we should call that out and we should say that is a problem and the recommendation is that that should somewhat be fixed.

CHAOYI LU

Okay, thanks, Peter. Are there any comments? Oh, okay, Jothan, please.

JOTHAN FRAKES

I almost clicked on the Zoom doing the vocals. Thank you, Jothan Frakes, for the record. So, I'm speaking about principles on the DNSSEC. I do have in my day job some perspective of being a registrar or understanding registrar inputs on this, and what I've discovered in similar work to Peter's good research is that there are many who do offer it through a control panel, but there's others where it's obfuscated or you have to make a special customer service request. They will fulfill the obligation that they have under the registrar accreditation agreement to add records, but they don't really make it convenient for their customers.

The other comment that I wanted to say here is that some of the things like Web3 domains or wallets or those types of things where you're doing authentication using text records often create a commercial interest for customers to ask for or demand DNSSEC, and it's actually been increasing adoption because registrars have historically dragged their feet about this because it wasn't as

commercially attractive as other things it could do. They didn't see a way to monetize it necessarily.

Now, as they're receiving, they're able to say, yes, we have it right here in the control panel. We can add those other features, or they're doing things with such name systems where attaching a text record that's in a DNSSEC-signed zone is offering commercial attraction, and that's what's kind of driving more and more adoption. However we might identify them. If we do identify registrars, it's probably good to say not that they don't support it, but that they support it by customer service requests or some maybe diminished thing than an active control panel. And we can signal that those registrars that do offer it are more attractive to work with or somehow in the documentation. Thank you.

PETER THOMASSEN: So just to clarify, it's all correct what you said. There are some that offer it through support tickets and all of that, but I was referring to cases where all requests just get flat out rejected. Hetzner, for example, does that, and we should call out that's a problem and people need to be compliant.

CHAOYI LU: Okay, so regarding the level and format of registrar support, Peter, do we have any more recent investigation on different registrars about how they support these features? Because the most recent

one that I can recall is from five years ago, so I wonder do we have other more recent data to support what we'll be claiming here?

PETER THOMASSEN: I keep an informal list for myself based on complaints from users of our DNSSEC hosting platform who failed to provision DS records for that reason with a registrar. I'll have to dig that up and you can share it next time we talk in our work party meeting. As for registrars who do offer that behind some sort of obstacle like a support ticket, I am not aware of any compilation of that.

CHAOYI LU: Okay, thank you. We look forward to that. Okay, Rick?

RICK WILHELM: Thank you, Chair. Rick Wilhelm. Regarding this, I certainly would not be disputing Peter's list or the folks within Peter's organization that contribute to the list, and I echo from my prior experience with Jotham's point about obstacles and not being offered the control panels, so all of that stuff can be true.

In my opinion, I don't think that an SSAC work product is the place to be basically filing a compliance complaint or a compliance report. And I think that if any member or organization of SSAC has an issue with any contracted party having issues that are contract compliance. They should just go file a compliance report with compliance and not wait for or mix in, wait for the production of an SSAC document or mix in their compliance concerns valid, which I

would presume that Peter's are, because I don't doubt the efficacy of his research or experience, certainly on this topic.

And we shouldn't be muddling up an SSAC document, because SSAC gives advice. But I don't think it's SSAC's job to be, like, trying to low-key bring in compliance by its work product, because we make findings and recommendations to the board and they're not compliance actions. So I would just suggest to Peter that you guys should just file a compliance complaint. And today there was a compliance action filed against a registrar that I'm sure that observant folks have noticed, not regarding DNSSEC, but about other things. And so I would just go ahead and do that. You don't need to wait for the paper. Thank you.

DAN GLUCK We have Billy.

VASYL "BILLY" BRATCHENKO Hi, guys. Did you do any investigation in terms of how registrars actually educate their customers about DNSSEC, or how they promote it, or do they focus only on being compliant?

CHAOYI LU Well, I believe we have distributed a short survey to the CPH group, and we are currently waiting for their responses about what measures have they done to promote DNSSEC, to educate DNSSEC

to their customers, so once we have the result, we'll put them in the report. Okay. Jim, next?

JAMES GALVIN In thinking about what Peter is saying, I think that we just need to be careful about how we approach what we want to say. I have, for more than 10 years, been saying that the contractual requirements about DNSSEC are inadequate, but not really having a way to capture that and make that point come forward, so Rick is right.

It's not our place to make a compliance complaint, and we need to be really careful about suggesting that that's what we're doing, because to make a compliance complaint, you've got to be party to the problem, so we're not really trying to register a domain name and try to add DNSSEC. We picked a registrar, and they don't have the right services, and it is known there are registrars who, the three cases that Peter identified, I mean, this is just well-known stuff. They either do it, they only do it through a special customer service ticket, or they just deny it and don't offer the services and say, sorry, we can't help you, even though they are required to do that.

So the question here is, can we observe in what we're doing here, as we are trying to identify issues that exist that are hindering the deployment of DNSSEC? And it is fair to observe that there are requirements that you have to do it. And if we can somehow reference or present, even if we anonymize, I don't think we should be naming people in our report. But if there's a way to capture the

evidence that here is something which is not happening, and we only have to state it as just another one of those things which is sort of in the way of deployment, don't capture it and state it as a compliance action or in some way anything like that.

It's just important to do that. You know, let's just stick to our technical side of things and provide the evidence and the data. And I think that we can do that and not make reference to contracts. Or we've got to find a way to make reference to the fact that this obligation is not present without making it a compliance or a legal issue. So wordsmithing will happen. Thanks.

CHAOYI LU Okay, thanks, Jim. Peter?

PETER THOMASSEN Yeah, so I agree with Rick, actually, that an DNSSEC report is not the place to make a general compliance complaint, but we're describing the state of things, right, and we do have texts on registrar support, and it's even on this slide, the first point. So, I don't think we should be ignoring the lack of DS provisioning support, even manual support, when we describe the state. And it doesn't mean that we have to name anyone, but we should reflect on the state of the landscape.

And you're right, I don't need to wait for the paper to be published. I can just file a compliance complaint for a particular registrar. In fact, FYI, I have done that. And it's a few years back, and receipt

was acknowledged, and nothing else happened. And I think that is not a good state of things. I'm not saying that this report should fix this, and I agree we have to be very careful what we say, but I don't think we should just ignore that.

CHAOYI LU Okay, thanks, Peter. I believe Maarten is next.

MAARTEN AERTSEN Hello, this is Maarten. So, before coming here, I reviewed what the work party was doing, and I noticed this really nice table you were making, listing the key findings and recommendations, which then turned into this slide. And one of the rows I really liked in the table was row eight, which was about post-signing validation.

There were some comments on the row, and one comment that stood out to me was by Warren asking with respect to the recommendation made on that row if it wasn't too like motherhood or apple pie or I'm not that well versed in American proverbs, but it's not on the slide and was just wondering if that's related to the comment. Because I think it seemed like an interesting observation in terms of finding. It seemed like a good recommendation to make and I wonder if it's now somehow gone or it's just a matter of a slide can only have so many words, and so I guess this is a question coming from ignorance and I'd be interested to hear more.

CHAOYI LU

No, it is not gone. It's still there. It's just we didn't have the space to put it up here. Yeah, so, a quick response to what you have about the Role 8 finding. I think there, , in one of the previous calls, we have agreed on that this, we need to be more specific on what has been the major causes of DNS outages from the historical events. And we will continue working on Section 3.2, and maybe bring up a summary of major causes, and make recommendations to specifically eliminate these causes. So, I don't think Role 8 is in its final form, but it is still there. Okay, so, Warren?

WARREN KUMARI

Yeah, Warren Kamari. I'm surprisingly going to have a bit of a rant. So, as Peter noted, there are a bunch of registrars who don't let you actually upload DS records, and, like, maybe this isn't the right place to name and shame, but at some point we have to be like, what the hell do we do about these sort of issues, right? It's a compliance issue. Registrars are required to do it. Peter has reported a number of ones who don't, many years ago. One, I have reported at least three many years ago, as far as I know, nothing has changed.

So, at what point do we go from, like, hey, it would be real nice if people actually, like, followed the RRA and did the thing, to, like, the shit's broken, yo. Like, this isn't working, being like, you have to follow the RRA, and then pointing out people don't, and nothing changes. Like, where do we go from here, right? Like, if compliance isn't going to actually enforce stuff, or if ICANN as an org isn't going

to enforce that people actually follow the RRA, kind of, why do we bother having it, and why do we bother submitting issues, right?

Like, your, when you reported, it was quite a while back, wasn't it? It was a couple of years? Yep. Is it fixed? Okay. I mean, is it worth your time to submit reports saying this? Like, I think a number of people have just given up submitting compliance reports, because, and yeah, as I, as I noted, like, this is going to be a bit of a rant, or was a bit of a rant, but it feels like one of the reasons DNSSEC isn't being deployed is because we have issues like this.

VASYL "BILLY" BRATCHENKO Can I?

DAN GLUCK I have Jothan, but Billy looks like he's about to jump out of his seat, so we'll see what Jothan has to say.

VASYL "BILLY" BRATCHENKO Okay. So, when I'm thinking about promoting by registrar, my first thought is that the benefits of DNSSEC are not even for all the users, so it should be targeted, I think, probably to, like, more email, big email providers, or when you become a big website, because if it's a personal blog, you most likely, you do not have technical capability to have properly configured DNSSEC, and that's why I would focus efforts on promoting this to those who actually can benefit from it. What do you think, guys?

CHAOYI LU: So, who, who will respond to that question? Warren, are you? Yeah, I'll let you jump in first.

WARREN KUMARI: I mean, yes, I kind of agree with you, but also not. I think for most personal users, at the moment, they're getting their service when they register a domain name, so, like, their registrar provides their DNS service for them, or they've outsourced it to desec.io. And so, like if your registrar provides it for you, if you look at the large organizations, it's only 7% of the top 100 have decided to sign, and it's not that it's too hard or complicated for them, it's they have chosen not to.

So, I don't think we can actually say if you're a really big org, you need to do this, because they've done the risk eval and decided they're not, and I don't think we can say, if you're tiny, you shouldn't do it. I think it's everybody has to make a risk assessment for themselves, and for some set of people, even if they are tiny, it's not their problem, someone else has done it for them, but I think I might still be on my rant mode, so.

DANIELLE RUTHERFORD: Just a heads up, we've got about two and a half minutes left for this work, Barry.

CHAOYI LU: Okay, so we still have three hands in the queue, and we'll probably get to that end on time, so Jothan, next.

JOTHAN FRAKES: I'll go very quickly, and it's my second bite of the apple, so that's the right way to do it. So with perspective as a registrar, and speaking more from that perspective to help color the conversation, I'm more principally thinking yes, it's identified in the registrar accreditation agreement.

I've identified that there's a process when ICANN enrolls a new registrar, or they have registrars renew, which includes, it's called an RIS, or Registrar Information Spreadsheet, or something, I forget what the specific term is, but it's an RIS, RIC, whatever, but if you had something where you fill out what your support of DNSSEC is, does someone have to email you, is it API based, is it a website, and then provide the URI of that, that would be information collected by ICANN, and it might kind of promote maybe a little more attention.

We're still in a situation where I think commercially registrars say, hey if I have attraction, if my customers are demanding this, then I'm going to implement it, otherwise I'm not, but anyway, there's also the registrar list. If you go to internet.net, you'll see there's a link, you can look at all the ICANN accredited registrars, and perhaps having ICANN list those that have support for DNSSEC, that might be a great way to, which should theoretically be all of them,

might be a good way to help promote the stragglers. Just some ideas that we might be able to do with ICANN's help. Thank you.

CHAOYI LU: Okay, thanks. Next one, Sourena.

SOURENA MAROOFI: Yes, my question is that mostly related to the complaint, so I want to know mostly from Jim, probably who knows more. So is it forbidden to do the complaint from a SSAC or name those registrars, or is it bad for the reputation, because at some point the SEC does some research, for example, come up with some results that there are these registrars not following the rules, so is it forbidden, or is it bad? I want to know. Thank you.

CHAOYI LU: Okay, Jim, you want to take that? You're also in the queue.

JAMES GALVIN: Yeah, so to answer your question, is it forbidden? Forbidden is a little too strong a word. I think what's important is you have to be party to the complaint. So, we can't just say a compliance thing without actually having SSAC taking it on as a body to participate and see that we did it, and then we want to complain that we didn't get what we wanted. Okay, that's a slightly different thing.

On the other side of it, I think it would be bad for us to do it, because the comment that I wanted to make is to restate what I was saying

before a little bit differently, at least speaking for myself personally, it's real important, in my mind essential, that we stick to speaking about technical subjects, and we should draw conclusions based on technical data, so that's the critical part in all of this. I am so sympathetic to where you are, Warren. You're 100% right. I'm on board with you and your page, but SSAC is not the right way to complain about whether or not compliance does its job. Okay, I'm really sensitive to that, too.

CHAOYI LU: Thanks. Okay, so to quickly summarize your point, you're saying that we can say that some anonymized registrars are not doing it correct, but we are not going to say that this is a compliance problem, right? Okay, we'll notice that in our report. So I think our time is up, so let's move on to the next work party.

DANIELLE RUTHERFORD: All righty, I think next on deck we have the DNS transparency work party, so I'll hand it over to Gautam and Raffaele.

GAUTAM AKIWATE: Thank you, Danielle. So basically what we figured was, since this is still a relatively new work party, we'd go over some background as to sort of motivate the DNS transparency work party, sort of recap some of the things that we have sort of settled on, and then talk a little bit about how we are viewing this might sort of come about, and sort of get folks' thoughts on it.

So for folks who have not been following DNS transparency work party sort of very closely, one of the motivating things that started the idea behind DNS transparency was these hijacks that targeted the DNS infrastructure of organizations, and in this case, they were targeting the DNS infrastructure that was not controlled by the organizations themselves, but at the registry and registrars, and I'll sort of walk through what that actually looked like.

And some of you have already seen this, but for making sure that we are all on the same page, sort of walking through it quickly. If you want to, let's say you're trying to log into your company's secure portal, you're going to enter that domain name on your browser. And then we are all familiar with how this recursive resolution sort of works.

And the thing that is interesting here is what happens if an organization wants to update its name service. So if it wants to update its name service, typically the way this works out is that the registrant uses, sort of communicates with the registrar that it wants to update name service, and then the registrar communicates that to the registry, which then updates the top level domain authoritative name service with the right name service. And so that's how it works.

And what essentially happened in these attacks, in the Sea Turtle attacks, is that the threat actors, in this case nation states, and from there on out other threat actors, basically targeted the registrars and the registries, which meant that they pretended that the

domain owner, the organization, wanted to change their name service. And sort of were like, hey registry, the domain owner wants their name service to change to ns1 and ns2evil.com, and please go forward and use that. And what that meant was that the users were then redirected to malicious servers. And one of the concerning things about this was like it bypassed traditional defenses like the TLS security also.

So this is where we are at. And the key insight that started a lot of this work was that the DNS configuration changes, and the domain owner has no insight behind what has changed and why it has changed. So the key idea is that we are going to ask folks to sort of log all of the DNS configuration changes at each of these levels.

The registry logs all of the changes, the registrar logs all of the changes, and when you sort of log all of this data, ideally, if you sort of do it in a near real time basis, there is a small chance that you can mitigate these security issues from happening. Or alternatively, there is a good chance that you can sort of audit through these logs. Did something bad happen to my domain? Which is something that we are currently unable to do.

And it turns out when you start collecting this data, there are a bunch of other use cases that also crop up. So if you're sort of doing this, I think Rod had pointed out that you could use this in the context of DNS abuse, where you could see domains sort of fast fluxing through a bunch of infrastructure, and that could be a key to uncover what is the infrastructure, what's happening there. And

then there are some of the other things that can also get flagged. For instance, certificate authorities could use this as a signal for, hey, this domain just changed ownership, maybe we shouldn't give it a TLS certificate immediately until it sort of matures a little bit. So that's at a high level what the key idea behind what we are trying to do.

And one of the things that the board party has discussed quite a bit, and I think have landed roughly in rough consensus on, is -- Warren seems really, really annoyed by -- okay, I was like I thought we agreed on the data types, what happened, but it turns out Warren is just thinking about something else, how dare you, but never mind. So roughly speaking, the data that we would think would be useful for logging at each of these three levels, and this is not saying that all of this data would be available at each of these levels, but the expectation is that whatever data they have, they can log.

So the first category of data is the zone data. So this would be the NS records, DS records, if they have them. This would be probably the most important set of data, and probably we should definitely have this data. Then what we are calling the category two data, which is some of the registration information. When was this domain registered, what happened, were there any changes, what are some of the EPP states, and sort of seeing if those sort of changed at any point of time. So if, for instance, a domain was under registrar lock, did that registrar lock go away, come back, sort of also could be a potential sign for some abuse.

And finally, the category three data, which is the registrant indicators, which is probably one of the more sensitive things that need to be discussed. And so, shorter to Warren, who's going to do a talk on this in his lightning talk, which is going to talk about how to do one mechanism where you can share registrant information in a privacy preserving way, potentially privacy preserving way across registrants.

So that is something that we are sort of aware of, like there are some privacy risks with sharing pseudo anonymized data with registrants. But we are not sort of making a valid determination of like should you publish this or not, but more along the lines this is like a comprehensive set of data that would be useful if folks were to log it, and this is not saying like you should log it.

So now that we have this data and these data types, the thing that we would like some input on and the thing that we have been sort of mulling over the last week is sort of thinking through how do these data collection models work. And the centralized model, which is something that Tim April and Warren Kumari had sort of previously proposed as part of their DNS transparency work project, was having this central aggregation hub where the central entity is sort of responsible for communicating with these different data contributors and sort of aggregating all of the data and then sort of having consumers that can sort of take it from this central aggregation hub. And like maybe a good way to sort of think about

this would be what CZDS is today, where there is a central place that everyone can sort of log on and get data from.

And then there is the independent publication model, which is also sort of possible, which is every entity essentially, like every registry registrar are doing their own logs. And this is not ideal because now there is some additional data processing that needs to happen to sort of deduplicate, to sort of figure out what's happening across. But that doesn't mean that this data is not valuable.

And what I would sort of akin this to is before CZDS was a thing, you had to go for zone files, you had to go to individual registry operators and had to sign individual contracts and then there emerged like a secondary market which sort of collected all of this data and presented it. And what then happened as the value of this data was sort of proven out, the centralized ecosystem sort of came up as a result of the value of the data.

So I guess where we are now is somewhere in the middle between centralized and hybrid. But we wanted to call out these two models and get your thoughts on does this make sense, because what we are sort of thinking about is our document needs to cover both potential models. Because this is like looking at the zone file evolution, how we get zone file evolution, like even though we have CZDS for ccTLDs, we are sort of still dependent on the independent publication model. So we are going to land somewhere in between central and independent publication. So we need to cover both.

And now I'm going to hand over to Raffaele, who's going to lead us through our discussion of this, and if folks have thoughts, please speak up.

RAFFAELE SOMMESE

Thanks, Gautam, for introducing this. So the discussion that we would like to have today is about the difference between these two models, the pros and cons, but not the pros and cons as technological pros and cons, but more as how these affect basically the actionability of the data that we can get out of the system if we, for example, we go from one or the other models in the way that we can describe them in the document, or also how this will affect the adoption of such a system.

For example, if an independent publication model is easier to adopt at the initial stage, like it was for CZDS at the beginning, before the system existed, where we are like independent parties that were able to publish the zone. I think it's worth discussing in the document.

And the second point, I mean it's correlated to the first, is also should the system provide only live data, so only stream of live changes, or should it also take care of historical data. And I guess this depends a lot on which kind of information we want to collect and how, which is the size of this information. And the last point we want to discuss is that we have been ongoing discussion about the fact, for example, there should be the possibility of I subscribe

just to the changes of my domain, the domain that I own in this system.

This requires, for example, filtering at server side. It may be expensive. So we'll also discuss whether operators should support this kind of filtering server side, or this should be like directly to the users that they can subscribe to the entire feed, or end up with a situation where there is like a secondary market that provides these filtering services.

So I'll open like the discussion with the first point maybe and see if folks have comments.

DAN GLUCK Okay, we have a hand from Maarten.

MAARTEN AERTSEN Yeah, so with respect to centralized or decentralized, I would ask what makes it more likely that you change the world with this work? I know you've been talking to at least one registry, so if one of these models has a more, like, a higher chance of being actually adopted, I would invest my time there, but that doesn't preclude from discussing the other one, I guess.

GAUTAM AKIWATE So I think the way we are sort of viewing this is an evolution. So I think both of these models are some things that we should cover. So I don't think this is a question of either or. I think both of these

models should exist. And the conversation is more along the lines of can we sort of think through what are some of the pros and cons for each of these approaches. And what are things that we should be sort of thinking of as we sort of propose these two approaches.

And I think what Raffaele also said was we see the independent publication model as something that we will probably start out with, as sort of registries are sort of testing out. And as we prove the value of the data, and as the value of the data gets proven, I think we will have more ammunition to go after the bigger fish, as to like making this something that all TLDs would like to participate in.

But I think the independent publication model is probably what we will have to begin with. And for instance, Raffaele has been in conversation with .se who have shown interest in doing an independent publication model. And so this is something that we would support. But then also we would, from Warren's experience, I think we would like to go to this world where we have a centralized entity, which will give us the most benefit. So the most benefit would be in the centralized world. But we are probably going to start out in the independent world. But we should figure out what this path looks like from independent to the centralized world. So if that makes sense. Raffaele, do you have anything to add?

RAFFAELE SOMMESE Yeah, I don't. We have Warren, did you…?

WARREN KUMARI

Yes, I mean, the centralized model is wildly useful, right? People have a centralized place they can go to to get the data. Everybody provides all the data there. The centralized system does all the deduplication, and independent model is, like, much, much, much, much, much, much less useful.

However, the only way you can have the centralized system is if ICANN mandates that it exists and mandates that everybody provides all their data to it. And I'm not going to say that's never going to happen, but that's basically never going to happen. So, like, this is a stepping stone, potentially, towards us being able to demonstrate that this is useful and then build something else.

RAFFAELE SOMMESE

And that basically was our comment on the fact that CZDS stemmed from effort of single operators in the past providing access to zone data, and then it became basically a centralized model after a while. We have Rick in the queue.

RICK WILHELM

Rick Wilhelm, for the transcript. Echoing Warren's point, but try and pick on some, there are some different points. I don't think that the centralized model is -- I don't think it's fair to compare it to CZDS because the data that's being contemplated is far more sensitive than is present in CZDS, both its content and its timeliness. It's been presented, and known as disputed, the fact

that there have been negative security outcomes for end users related to past abuses in this area that have involved this exact kind of data. And that was in an environment that was occurred years ago that is a far less hostile landscape than the current world in which we operate.

I think we all agree that the world is growing more aggressive and things are only moving faster. And when the VeriSign service was shut down, that was a world in which, by comparison, looks a fair bit simpler to the world in which we operate now. So it's only gotten more hazardous out there.

I would offer that while this work product may contemplate and examine these other options, as far as a recommendation, we would be better off talking about standard publication models that facilitate an easier interchange between those that are willing to publish their data and those that want to subscribe to that sort of data, so that there's easier interoperability for researchers with those that want to publish and that sort of thing. But we really got to be concerned about, since we're trying to improve security, creating an attack surface that is that has a potential of making things a lot worse.

And when I see a model that has a big blob of data in the middle, that seems to be creating an attack surface to me that has a potential of doing way more harm than the kind of security research that we're talking about. Thank you.

GAUTAM AKIWATE So I think that makes sense. And the VeriSign Rapid Zone updates was 20 years ago, and I think the world has moved quite a lot. And you're right, it is a lot more aggressive. But at the same time, a lot of this data already exists out there. And in terms of certificate transparency logs, they're already out there. All of this data is like there are companies that do newly observed domain. So I don't think it is fair to say that this data doesn't exist out there. The attackers have access to a lot of this data. What actually is missing is for organizations to be able to access this data.

So I think you're right in that we need to think through what it means for us, for a central place to have all of this data, and make sure that we have the appropriate checks in place. But then also the reality is like a lot of this data is already public and actually accessible to anyone, like in terms of CT logs. And a lot of Raffaele's research actually shows that you can recreate a lot of newly observed domains just by looking at CT logs. And 70% of domains actually have a TLS certificate. So you essentially have a data source that already exists that is used by attackers. So what additional threat vector are we adding is something that we should actually consider. But I agree that things are.

GAUTAM AKIWATE Jim?

JAMES GALVIN

So yeah, Jim Galvin for the record. I want to build on Rick's point and come back to a question. I'm sorry, I forget who asked it over here on the side, commented about, oh, I think it was you, Maarten, about whether, which one of these two choices of models is the right thing. I think it's important for SSAC to stay on, to not solution solve unless we've got a very compelling documented SSR reason for picking one or the other.

Okay. So it's important for us to offer options and then offer the advantages and disadvantages of doing those things. Because that's not the way -- to do it any other way doesn't really get it in front of the contracted parties who you want to do this. If you want registries to take this on, we then have to take it to them and get them to agree to do it. Or you need some kind of PDP and create a consensus policy that makes it happen. So there's a process for getting that. We should simply lay out what's needed and why. And I think that's really important.

And to this privacy question, it's a similar kind of concern. Whether or not something is private information or not is a legal question, not a technical one. And in spite of the fact that we have lawyers as part of our SSAC team here, and I'm sure that they have an opinion and I'd love to hear it and if we can capture some of that that's great. However, our legal review of anything is going to pale in comparison to the lawyers who sit with the registries and registrars who are going to look at this and make a decision.

So it's better that we should talk about why from an SSR point of view something is good or bad. And in fact we should compare it to being good or bad against CZDS. Technical point of view I think that's a valid thing for us to do. And then the next step after our document is published and acknowledged by the Board is to work with the community to see what we can do to advance this forward if we want people to do it. That's the right way to do this. So I just want to be careful about all of that. Thanks.

DANIELLE RUTHERFORD: Real quick, we've got a time check, five minutes left for this work party session, for this work party in this session.

GAUTAM AKIWATE: Billy, go for it.

VASYL "BILLY" BRATCHENKO: Can we achieve our goal without sharing the entire information to the world? I mean, probably it could be targeted with some authorized like investigators and so on without providing access to everybody. And moreover, it can be traceable back who actually requested this or that data. So if it's more like a closed model, probably it is more achievable. What do you think?

GAUTAM AKIWATE: So I think that's primarily the reason we are sort of pitching the centralized versus independent, where we suspect that in the

beginning, folks are going to have concerns about like the data. And I think Rick brought up like excellent points about like there is some sensitivity around this data and we don't know how attackers might sort of exploit it. So this is exactly what we are hoping will encourage folks to adopt it.

Ideally, we would like this data to be public eventually. But I think in the beginning, I think it's perfectly fine for folks to be like, I think this data is sensitive. Let me sort of trial this by having a smaller set of folks like do this. And then we can sort of flesh out like do the lessons from 20 years ago still hold true today, which is so the various and rapid zone update data got abused. Does that still hold true today or are we seeing the benefits outweighing the risks?

And in this case, I hope that the answer is like the additional data of information that we are giving attackers is not much, but the auditing and mitigation capabilities that we get as a result are significantly higher from this. So to answer your question, yes, I think the independent publication model is exactly why we are discussing this. That said, we would eventually like to see this data to be public.

VASYL "BILLY" BRATCHENKO Why?

GAUTAM AKIWATE So I think we can talk a lot. Like I think there are thoughts as to what would be the benefits of having this data to be public. I think

it's easier to access. And what we have found is that with the certificate transparency log data, like anyone being able to audit has meant that if the organization themselves are not keeping up to date as to what's happening, there are other entities that can fill in the role for you and like look at trends that are broader than like things that are targeting a specific organization.

And yes, like some of this role could be filled by security companies, but more generally like researchers, like academics could look at this data and use it for things that uncover deeper problems. So like as a personal example, the reason CCDS data is not technically public, but it is effectively public because once people collate it, a lot of this data is public, you can sort of get it from a bunch of people. And what that allowed us to do, and the reason I got involved with SSAC was being able to find operational risks that were hidden. And the only way we could do that was because we had access to that data.

VASYL "BILLY" BRATCHENKO But we can authorize academic researchers and trade them back in case.

GAUTAM AKIWATE Sure, which is why the independent publication model -- this is a personal preference, which is we would like the data to be public, but we understand that folks have different risk appetite and we want to accommodate that.

RAFFAELE SOMMESE

I would like to add to that, that independent publication models allow the different registry and registrar to vet whatever the request that comes in is deemed to be valuable and it's deemed to be worth the access to the data. And that's outsource of the discussion towards then a centralized model in which all these instance are accommodated basically. But again, it's an evolving process and I think both models should be considered exactly for this.

RICK WILHELM

I'll be very quick. There was some question or some statement like, we're not sure about how these things are going to be, this data could be abused, which kind of ties in a little bit to the thread that Billy was pulling on. Sorry, this is Rick Wilhelm for the transcript.

I just wanted to bring up something that Matt had said a couple of months ago, April 15th, in his note to the list. Domain life cycle events are not just neutral signals, they are high value triggers. A neutral, newly registered domain is in a particularly fragile state. The registrant is often in the middle of setup, configuring stuff, and they are more likely to trust communications that appear to be part of that process. If you expose that event stream with very low latency, you're effectively giving adversaries a playbook for when to engage. It becomes trivial to draft onboarding themed phishing

emails, impersonate registrars or hosting providers or exploit configuration gaps before the domain is fully secured.

That's why this is a risk. We're by exposing that event stream, we're exposing new registrants, people who are configuring new domains, to this kind of risk. To steal a term from privacy legislation, regulation, it's a balancing test. I think we need to be extraordinarily careful to be advocating for putting out more data that puts people at risk for these sorts of things. Thank you.

GAUTAM AKIWATE

I think that's fair, except all of this data is already public in form of certificate transparency logs. The same arguments apply for certificate transparency logs where they are in a sensitive space. When you get a certificate transparency, like a TLS certificate, you have a web server that is in the process of being set up and doesn't have a TLS certificate yet. You could be subject to attack.

I think a lot of these things have been true and a lot of this data is already public. I agree with Matt that the registration events are sensitive. The question is, what is the data that is already not out there? What is the additional data that we would expose people to because of this data as opposed to things that people can put together based on certificate transparency already?

I think the argument that we are having is, in the last 20 years, there have been additional data sources, public data sources that have added that already make this data public. DNS is not the only

source of information as a signal for, oh, something has been registered.

For instance, immediately issues a TLS certificate whenever a domain gets registered with them. That is a signal. That already sort of tips the hand to attackers that, oh, a domain has been registered. Are we saying that? I think that is the thing that we need to be careful of is, what is the additional risk that we are imposing as a result of this data? I think that is going to be an ongoing conversation, and we need some time for SAC126 discussion. We will close with that, and we will continue discussing this in the next work party meeting. Thank you.

DANIELLE RUTHERFORD: All right. The final agenda for this SSAC work session is a discussion on the SAC126 update proposed by Peter Thomassen. Peter, go ahead.

PETER THOMASSEN: Thank you. Next slide, please. In the SAC126 report on DS automation, we had given recommendations, three of them. One was that if a registry or registrar wants to do DS automation, then the current recommended interoperable mechanism for that would be CDN, CDS, CDN as key type of DS update requests specified in a bunch of RFCs.

Second recommendation was that ICANN should support entities who want to do that, such as by facilitating a fast track RCEP. The

third recommendation was that further operational guidance should be developed around implementation aspects where different players in the past have made different choices, like, for example, with respect to registration locks and DS automation.

This is what we did two years ago. And since then, a bunch of things have changed. And today, we're talking about how or whether to align our advice with the new developments. I'll focus on the section 4.4 that's listed here in recommendation three, because that was the starting point of the developments that happened after. Next slide.

Okay. So, the section 4.4 lists a number of operational considerations where implementers of DS automation have to make practical decisions, like, for example, who should do it, the registrar or the registry, what are the acceptance checks, should you do that once or should you do that continuously, what's going on with TTLs for quick rollback, and so on and so forth. I would assume that all of you have memorized this two years ago, so I'll not read out of it.

But in appendix B of the document, we didn't give any recommendations about these points, but instead, we gave starting points and ideas and suggestions for further discussion of these things so that somebody could develop guidance about it. And as you recall from the previous slide, recommendation three was actually to come up with such guidance. And as a result of that, there is now an IETF document that settles these questions based

on what the SSAC started thinking about in that appendix. So that is the background context. Next slide, please.

So since we published this, the IETF has approved actually last month, the draft IETF DNSOP-DS Automation, we're currently waiting for an RFC number to be assigned. So this was co-edited by Steve Sheng, who at the time was at ICANN and was also overseeing the SAC126 report. It's a BCP type of document and it's titled operational recommendations for DNSSEC delegation signer automation.

So the other major development is now that ICANN is working on recommendation 2 of our report, which is facilitating and supporting registries and registrars who want to implement that type of stuff, for example, by enabling a fast track RSEP. And in that context, they are doing a security and stability assessment. And while they did that, they pointed out that there are differences between what the IETF now has approved and between what's in our appendix B, which was the starting point for the IETF document. And now it seems like some clarification is needed to move this further along about what's authoritative and whether the SSAC thinks that the actual way of doing this is appendix B or the IETF document.

And yeah, so this is why we're here. Some other news since then are that there are two other RFCs, which have been published since then. The first one of them is RFC 9859, which is about the generalized DNS notifications where a child operator can ping a

parent and tell them, please look for a DS update request now. And the other one, 9975, is about consistency requirements between CDS and CDNS key records if you publish both. Both of these topics are mentioned in our report. And at the time, they were either suggested to be done or they were under development and labeled as in progress in our report.

So these are just two things that have also seen updates in the real world after our report was published. But the main reason we're talking is the first one, and ICANN org's security and stability assessment. I will go now further into the two things where differences were pointed out. Next slide, please.

So on the left-hand side, we have what our appendix B suggests. On the right-hand side, we have what the IETF document suggests. The first topic is about the interaction of DS automation and registration locks, particularly when there is a server update prohibited EPP status, which is sometimes called a registry lock. And the question is then when that lock is in place, does that mean that you are prohibited from doing DS automation or that you can continue to do it? In the SSAC report, the appendix B did say that or suggest that that maybe should be prohibited at that point.

The reason for that was mainly to have a conservative position to start from and see if there's any good arguments that can be made why that is not a problem, in which case the restriction could be loosened. And in fact, it turned out in the IETF standardization process with input from Scott Hollenbeck, who's an EPP expert,

that there is no contradiction between this EPP status and actually performing DS automation.

There's various reasons for that. So the EPP specification itself says that registry-side policy can actually change the status of the registration object independently of this record, just not based on an EPP transaction, but DS automation done by the registry isn't an EPP transaction. And yeah, so that was the main argument. I've linked Scott Hollenbeck's input on the mailing list here. And if anyone is interested, we can review it in detail or you can do it offline.

And yeah, so this was one point. And the other point was that the Appendix B in our report suggested that children publish both CDS and CDNS key records, information which are in a way redundant because one is more or less the hash of the other, and that parents should require that both are actually published and otherwise not act. The reason for this was that that dichotomy is suboptimal, let's say, and if we ever want to, I mean, we, not the SSAC, but the community ever wants to deprecate one of the different types, then it's important that everyone at that time can easily switch to the other type. And that is most easily achieved if everybody just uses both types now.

The IETF rejected that idea, which is fine. It's not a security or stability problem at all. So the IETF document now recommends that children still publish both of these things because that's most interoperable. They don't need to know what the parent prefers,

but the parent is advised in the IETF document to not abort the operation when only one of the things is present. Instead, they should see whether there is a contradiction between them if they are both present. And if there is contradiction, they should not apply the update. But as long as things are plausible, they should continue.

So these are the two changes. And based on this, I was thinking last week and proposing on the SSAC list that maybe we should issue a statement that the SSAC is of the opinion that the developments starting from Appendix B were reasonable in the IETF and it's all in the spirit of the report. And there is no reason to be concerned about what the IETF said and that we endorse essentially that the future process at the ICANN evaluation is based off the IETF BCP.

And yeah, then I realized, oh, these other RFCs also appeared in the talk to Ram and Danielle. And the idea was surfaced that maybe it is easier instead of explaining what the diff is to just actually copy the document and make minor edits and have a new document, the SEC126v2, that just has small edits and explains what the situation is now. So this is the proposal. It's on the next slide. And I already see Warren's hand, which is nice. I like your hand.

So the proposal is to issue that update to record the, sorry, to align the SSAC's advice with the community consensus from the IETF resulting those questions that I showed on the second slide. The two recommendations that would be given then are the ones displayed here. They are essentially the same as before, except

that the DS automation and notification RFCs are added. The second recommendation is unchanged. And the third one about developing this recommendation, sorry, but developing the additional guidance is gone because that's completed. So this is a minor, and then there would be editorial updates, like topical ones like this guidance should be developed, would be turned has been developed and general things like updating RFC numbers and stuff like that.

So I'd like to put this in front of the SSAC and see what people think. And also we can probably, if people are interested, actually put the draft v2 document on the projector, but I don't think we need to do it now. I think we can first have a general discussion and then dive into the details if folks feel like that. I don't know what the queue order is, so Warren and then Maarten.

WARREN KUMARI

Yeah, I mean, it all seems a little silly that we have to publish a whole new document to be like, by the way, we said like, we're going to develop this process and now there's thing done, right? Like, it feels like this is a symptom of a larger issue that in order to fix the thing we have to go through all of this process, that's really the only way that we can communicate to ICANN, please read what the words say. I guess we have to do that, but at some point we should discuss how we make this process less filled with process and slow.

PETER THOMASSEN Yeah, I had the same concern actually. And I thought, oh, if we do v2, we need a work party and a chart and all of that. I was told that's not needed. Apparently, probably Daniella or Ram --

WARREN KUMARI We don't need that. It's done.

PETER THOMASSEN I agree. It should be easier, but yeah.

DAN GLUCK Maarten?

MAARTEN AERTSEN I support solving the problem that's surfacing and let's do it in the simplest possible way. So if that means Peter does work with Danielle to propose some edits and we have a seven day window where people can express that they don't like it. Great. If it means sending out the document that I've already read that Peter sent to the mailing list, also great. I guess the relevant question to the group is, does anyone expect to object to any of these methods? Because then we will waste a bunch of time, which to Warren's point is not very valuable activity.

DAN GLUCK Go to Jim, and then Danielle.

JAMES GALVIN I'll just add for context here. I agree completely with Warren that this does seem like a little bit of overkill, but on the other hand, keep in mind that documentation is important to ICANN org as it does its thing. I'll simply observe that the Board actually had a particular issue with the delegation in a resolution recently, which is still resolving itself in the IETF. And once that gets resolved there, we actually had a bit of communication where we had to say, look, the Board will just do whatever the IETF wants. Just tell us what it is and then we'll do another resolution to fix it. So we get past it.

I'm with Maarten. As long as we don't have any objections to the technical substance, let's let staff figure out what the right way is to make this happen and just go with it and move along. Thanks.

DANIELLE RUTHERFORD Danielle Rutherford, ICANN Org, support staff for the SSAC for the record. I hope I'm not repeating what you said, Peter. I just want to clarify why the recommendation came out to do a V2 instead of the initial succinct draft that Peter put out. That was my recommendation because the original draft that Peter put out on the SSAC mailing list contained updated SSAC recommendations.

My recommendation then and remains now, if you are updating a document's SSAC enumerated recommendations, the text of that, in general, the operating ethos from the SSAC has been that most

of the SSAC NNN series reports should be as evergreen as possible. And in that sense, I think if you want to update a recommendation, I do recommend you go through a V2 process.

If the intent is to just add additional clarifying information, we do have a process to just send communications to ICANN org through the ARR process. Or you can publish an addendum similar to what Peter had already drafted. I just don't think that's the appropriate mechanism for updating previous recommendations.

PETER THOMASSEN: Yeah, so maybe that raises the question whether we should update the recommendations or just provide context. I think for maximum clarity, it's best, like, when we list RFC numbers in recommendation one to actually add the new one. So that would be an update. And that is why that's my slide preference, but it's certainly up for discussion.

DAN GLUCK: Raffaele, and then Warren.

RAFFAELE SOMMESE: Raffaele Sommese, for the record. I think we should go for the path of less resistance. I mean, whatever the IETF, you define the IETF that's technically sound, can be reproduced in the recommendation. If that's an easier process to do, we should go for that.

WARREN KUMARI Okay, I was waiting for you to say, Warren. See, I mean, Danielle's suggestion is reasonable. Just, like, hopefully us publishing a V2 document should just be like, here's the V2 document. It's exactly the same as the V2 document. We just listed another RFC and we're saying read the doc. And it doesn't end up being many, many weeks of going round in circle. Does anybody have an auth 48 objection?

BARRY LEIBA Well, I'll just say what the IETF does with this. The answer is nothing. If an update to the RFC comes out, they don't update all the documents that reference the old RFC. They just leave it and people will find it. The information that's there is sufficient.

DANIELLE RUTHERFORD Just to respond to you, Warren, and I was typing this out process-wise, anything written, an ARR response, an addendum V2 would all have to go through SSAC consensus. And opening things up for V2, I think in the past has been very precise. It's not opening up any other content in the document. It is specifically to align with the newly published or to be published RFC.

GAUTAM AKIWATE So we have done the V2 process before? Okay. So it seems like it's well understood what a V2 process looks like. And we have a mechanism for that. So I think that seems to be the most

reasonable thing here because it updates the recommendations to come in line, makes the recommendation make more sense when people go back and look at SAC126. So I think it sounds like a good idea.

WARREN KUMARI: I mean, yeah, I'm supportive of us doing the V2 thing, just like hopefully this doesn't turn into another one of the like, and then three weeks later, somebody is like, I don't like the color that you used and the font or something stupid.

JAMES GALVIN: Okay. So to bring this to a conclusion and state this precisely, so we have no objection to the technical substance of what's going on here and we're going to turn it over to staff to pick a process and follow it. The V2 document seems to be what's on the table from Danielle. We'll just do that and we'll work appropriate consensus on the list as the process requires. Anyone have any objections to that? And if not, I think we're good, right?

DANIELLE RUTHERFORD: Jothan had his hand raised, but it was before you asked for objections, so.

JOTHAN FRAKES: Yeah, it was not related to an objection. I support it.

PETER THOMASSEN Okay. So then I'd say that perhaps staff can prepare a red line version and then we send it across the list and otherwise call it lunch.

DANIELLE RUTHERFORD All righty. Thank you, everyone, for joining. You can please stop the recording.

[END OF TRANSCRIPTION]