ICANN82 | CF – SSAC Open: Source Software Work Party Meeting
Saturday, March 08, 2025 – 16:30 to 17:30 PST

JOHN EMERY: One minute warning, everyone, one minute warning. All right. We are at the scheduled start time. Can you please start the recording.

KATHY SCHNITT: Thank you. Hello and welcome to the Open-Source Software Work Party. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior and the ICANN Community Anti-Harassment Policy. And with that, I'm going to hand it over to John Emery.

JOHN EMERY: Thank you so much, Kathy. I'll go ahead and share the screen real quick just so we can see a little bit of the agenda for today. And Maarten, if you want to take us through what our plan is.

MAARTEN AERTSEN: Hello, everyone. I'm happy to see so many people attending. Today, we will follow up on some of the action items from last Friday. Some people volunteered text, and I just want to go through that. Then I have two items for discussion. One on Section 3 and one on the SSR subsections basically. We'll review a charter type line and there's two things we won't be doing today, which is I would like to invite many of you to think with me collectively about

some sketches for diagrams because we have some diagrams in a document but we could use another one for the introduction, and I think thinking about diagrams is way easier when you're in person than when you're online. So, it's not for now with mics but I would like to invite you to join some sketching maybe in the hallway or whatever.

And then the next normal meeting would be March 21st, which is not this Friday but next Friday. I just checked with John whether he would prefer to skip a week as is usual for us, but he actually is fine doing that meeting then, so we'll resume online on March 21st. And what's on the agenda for that meeting is to discuss the survey that ran during February that Vicky analyzed. And yeah, attend if you like to hear about that. So that's for the agenda for today, and I propose we get started.

So, there were four people that volunteered or were asked to provide text. We did some—and I don't think Geoff is here, is he? Is he online maybe? Okay. I was hoping he would be willing. I was hoping he would maybe give us an input. I'll try to see if I can talk to him later this week. So, Raffaelle gave some more input on the methodology for the appendix, which is very nice. I think Warren is finalizing his section on resolvers in the appendix. And then John was meaning to copy edit the appendix. John is not here either, is he? I think he's meant to arrive. He is here in person later this week?

KATHY SCHNITT He's not attending, and he sent an email earlier that he wasn't even going to be able to join remotely.

MAARTEN AERTSEN Oh, okay. Well, he offered to do some copy editing. I'm not sure if he's got to it. But that's the action items as we have them. I think we could go to the next slide.

So maybe a question before we dive into this part. How many of you have looked at the document recently? Because I see a lot of people that have attended some meetings, but maybe you have not been super active. So, I can give a little bit of an overview of the document if that's helpful. Please raise your hand if you would like to have a little bit of an overview to maybe get back into the—right? Yeah, great. So then if we can, can we switch from the slides to the doc? This is great.

So, this document has five sections basically. So, the first one is an introduction to try and get our policy audience to give a broad picture of the two aspects that we're talking about. The domain name registration side, and the DNS side which is like where queries are answered. Then there's a little introduction about open source and that's basically where we try and set the stage for the audience. I think this bit of the document is in pretty good shape. It could use your review even if you haven't done any work in this work party because it's intended to be accessible to a general audience. So, if you think about doing anything, this is a great place

to start. I think it was written by Robert and by myself, and I think that's it for the introduction.

So, then Section 2 dives into what the charter calls the landscape. It introduces this notion that both DNS and domain name infrastructure is built on open-source software. It's like the human readable version of the appendix where in the appendix we have some statistics about where is open-source software actually used. And here it kind of introduces the different bits that are part of DNS like authoritatives, recursives, and etc. And then it has some statistics that we gathered about what the role is of open source in this context.

So maybe one final note about the introduction that I should have made is that that's why we could use that diagram to kind of introduce these different sides. You can see in this section there are also already some diagrams, and it would be nice to have some in the beginning as well.

So, the other bit that is in here is the domain name registration infrastructure, which is like the EPP side of things if you will. I recently made some edits here because I think this is the part of the report where we don't have great data. So, for the DNS side, we have good statistics to offer our audience which are also new like not available elsewhere, which is I think what makes a publication valuable. For the EPP side of things, we don't have great data and there's some comments in the sideline where I think John Levine gathered some information, but I'm not sure if the way we summarize it here is accurate. So, this bit needs a bit of work.

So, then Section 3 moves from the landscape description of what the world is like towards a discussion of what the characteristics are of this software. So, this is basically a general overview of, so what is open source but more specifically like why is it different from proprietary software, again, for a policymaker audience. And then it is not like about open source generally but about the open-source software that's used in DNS and domain name registration.

So, this is a section that Vittorio did a lot of work on. And when we accommodated Geoff in the fall of last year, we tried to set up some meetings that were better for the Asia time zone, and Geoff then expressed an interest to contribute something to this section, which is why his name is on it. But right now, we don't have his perspective in here but I would like to get it because it's a bit different than the perspective that's currently in the text. So, I would like to kind of blend these perspectives, and I hope to talk to him later this week.

So, 3.1 is basically a summary of what we discussed with Geoff. It needs to like be integrated and then the subsection can go. The second subsection of this is still about characteristics but it's like the SSR implications. So, from this section onwards, we basically get into why SSAC exists because it discusses the security implications of open source being in in places. And I think this subsection is on our agenda for today because I think we can do a bit more work considering various angles and expanding this.

Moving on to Section 4 is where we move from characteristics of the software to operating open-source software. And this section is

where I think the document needs most work, because it's still very short, very scarce of any content. I think Ram promised to offer some content. He mentioned that he will do this, so that's great. And here again we have a description of operating open-source software and then also, maybe implications. And there's some bullets there, so it just needs to be expanded into text.

Then Section 5 is basically where we dive into the policy side of things, and I made a first draft of this second section, so I could use feedback. But it basically talks about assumptions that people have, which may not actually apply in the context of open-source software. That's the whole document. There's of course a conclusion section but that is something that we can discuss at a later stage. And then, the appendix A corresponds to Section 2. It's basically the data that we have about the landscape.

And appendix B is something that might or might not disappear in the end but it's the discussion of what the survey results were, and I will not spoil that today. If you want, you can have a look but we will dedicate the next work party meeting to discussing the survey results. And Vicky who couldn't be here today did this analysis work, and I'd like her to share that work. So maybe that's a high-level overview of the of the document as it is today, and maybe this is a good time to just look around and see how many people—if there's questions whether I lost you somewhere halfway, etc. I don't see that many confused faces, so I think we can just proceed. Again, I'm getting a smile, okay.

So then maybe we can move on to the specific topics, and let me open my local copy of the slides. So, I think what I would like to do is discuss bullet three on this slide. Hadia?

HADIA ELMINIAWI Could you please just put the link again in the chat for the document?

JOHN EMERY For the document? Yes.

MAARTEN AERTSEN Thank you, John. Since Geoff is not here, I think we'll skip one which was discussing the different perspectives on Section 3. I think what's useful to do now is to look at the subsections of Section 3 and 4 about the SSR concerns. So, if you scroll down about a page, we get to a heading called implications for security stability and reliability. And if you open the link that John just shared into the document, it would be helpful if you could read that bit of text, and maybe give feedback or brainstorm together like what are these implications that should be in a document.

So, I'll pause talking here and give you a minute to read. So that's Section 3.2, and the link is in the meeting chat. So, to give a little bit of a summary of what's written here, I just marked a bunch of words bold, which is basically a list of bullets what we have right now. So, if there is some talk about centralization, there is some benefits of transparency. There is a reference to protocol development being

like there is some talk about absence of funding. And I think that's what we have right now, so we are far from done.

Oh, I missed the diversity factor. So, we are a room of knowledgeable people, how about I hand the mic to the audience to see if we can come up with other angles into this topic. People still reading or maybe checking their email. I'm just looking around what the room is thinking, doing. I'm noticing that it kind of shows that the people that have provided most text are not here today like with Warren in Bangkok and Vicky not able to make it. So, it's kind of—I need you to give me some energy here.

JOHN EMERY: We know it's the end of the day, but you all have the skills. I guess, the question would be too, we have these four bullet points, diversity, centralization, transparency, and funding. Is there anything else that we should include in this section? Some important implications for SSR that we may be overlooking in this context. It's always easier to start from a bullet point, and then build on it or does this pretty much encapsulate everything that we're looking at? Yeah, Rod.

ROD RASMUSSEN: So, where would you factor in malicious code insertions and look-alikes? So, there's a very common problem these days with—they call it GitHub fishing for lack of a better term. Do you think that's well covered in these bullet points in maybe the transparency setting or…?

MAARTEN AERTSEN No, I think that could be an angle that we could add here. So, I guess, the argument would be that this is somehow simpler or I guess, you need to compare and contrast to other systems then.

ROD RASMUSSEN Right. Well, if you've got proprietary code, it's a lot harder to hack that, right, versus open source. The winner is whoever has the repository and it's really easy to get people to download libraries that are infected and or include them and things like that, so.

MAARTEN AERTSEN Right.

ROD RASMUSSEN I'm talking from the supply chain perspective.

MAARTEN AERTSEN Yeah. No, I understand. So, I guess, the question would be—I would like to have something like this. And then the question is, what specific aspect will we focus on? Because you can talk about, let's take a name server as an example. So, you have NSD, a name server that we develop at my workplace. And so, would the concern you want to describe be that as a drive by contributor, you have a chance to inject code in a release project, or would it be that software has the dependencies, and therefore, you could inject something through a dependency, or maybe something else like—

ROD RASMUSSEN

Two things that I would say. One is the dependencies, right, so libraries are the particular one. Especially when people are building new code. I want to include my latest AI whiz-bang things because I'm going to incorporate my DNS, a lot of filtering and things like that. So, I want to use AI to do that and so, I'm going to download some cool thing. That's really hard if you don't understand the technology itself, and include that library. It could be malicious, right? So, that's the supply chain aspect.

The other aspect is kind of what I call the phishing aspect, which is somebody wants your product, but somebody else has named something similar to it or has taken over an account like a GitHub account. So, there's provenance or assurance around that the software is actually safe. This goes into why open source is great because everybody can examine the code, but a lot of people just assume that since it's there, and it's been there for a long time or what have you, that it's good.

MAARTEN AERTSEN

Right. So, this would be like the distribution aspect. How do operators get the software that they're using. So, I would like to like add some bullets to cover these. I guess, with both of them, I would ask whether like the threat model you described would be unique to open source. For example, if you look at the dependency angle, we know that most proprietary software also includes open-source dependencies. So, in that sense, compromising a dependency

chain is not a unique thing for open-source software because it's quite rare for complex software not to include.

ROD RASMUSSEN Yeah, these days.

MAARTEN AERTSEN Yeah.

ROD RASMUSSEN Yeah. I guess, the difference there being, if you're working strictly in an open-source mode, you don't have potentially the resources of a proprietary system. If you're looking at Microsoft doing another DNS resolver, they've got a whole bunch of people doing validity checks and making sure that the bomb is properly ascribed versus somebody doing a home project, right? But I get what you're talking about, it's not necessarily unique because so many people have adopted open source into their own proprietary products.

But I think that the fishing angle is an interesting one especially given that you have two forms there. One is just a look alike from the naming perspective, and then the other being the, I can get access to your GitHub or whatever account that has that, and I can put my own code in there pretending it's you, the developer.

JOHN EMERY Yeah. So, we have Russ and then Robert, and I just also want to highlight the question from Farzaneh in the chat as well. So, Russ you go first.

RUSS MUNDY Thanks very much. One thing, it might or might not be worth mentioning in the document, but I can say with absolute certainty, there's a very long tradition in infrastructure operators and provisioning to make preferential use of open-source software in the infrastructure against the proprietary software. Now, if that is something that is worth noting in the document, I'm not sure, but it certainly has been the case for many years in the DNS realm of things. Certainly, in the operating side of it, is an operational preference to make use of open-source code. And that's, whether you'd call it momentum or not, I'm not sure. So, it's something to think about whether or not we want to include something about that. Thanks.

MAARTEN AERTSEN Thank you, Russ. So, if it's about the tendencies of operators to prefer this software, that would be—

RUSS MUNDY I don't hear you, Maarten.

MAARTEN AERTSEN Oh, I seem to be hearable in the room. So, one, two, one two. So, Russ, is audio good now? Otherwise, I'll come back to you in a bit. Maybe someone can tell Russ—

JOHN EMERY It looks like he can't hear.

MAARTEN AERTSEN Okay.

JOHN EMERY Russ, we'll come back to you in a minute. I'll type it in the chat. Robert?

ROBERT GUERRA Yeah, so just slightly different. Maybe you're talking kind of bullet points that we may want to add. So, I think on the transparency kind of open source, it might be kind of peer review to maybe get to what Rod was saying and stuff like that in terms of like code being inserted and things like that is just that when it works, the changes get in but people take a look at it before it gets adopted. And so, I think we may want to add something like that is that there are more eyes on it, and some of the more successful ones have a peer review process. And so, that is a value.

MAARTEN AERTSEN So, I think I do recall that this is already in a text somewhere because we also—yeah, it's in the transparency paragraph because

it also lists that the major open-source application in the DNS space have been the active subject of academic study for decades, and it continues to this day. So, two of my colleagues for example, get multiple reports from academic researchers each quarter more or less. And I don't think that is common in open source generally, but it is common in the DNS space, that there's this much interest. So, I like this comment, I think—

ROBERT GUERRA Sorry. It was just more—you're kind of talking kind of you have diversity, transparency, so just maybe bringing that further up, so people see it as a core value at the beginning and then you just describe it later. So even if it's there, I think just printing it in bold—

MAARTEN AERTSEN But we're not keeping the bold text when this is published. It's just to make this meeting a little bit simpler for people that haven't been reading a lot.

ROBERT GUERRA So, the one thing now is maybe to have that emphasis shown up earlier when people see it knowing that it's coming later. So, that's the suggestion.

MAARTEN AERTSEN Okay, thanks. Do we have audio or will we go to some—

JOHN EMERY Russ, do you hear us now? Give us a thumbs up.

RUSS MUNDY Got you. It was a problem here, sorry.

JOHN EMERY Okay. Maarten, do you want to answer Russ's question now?

MAARTEN AERTSEN Yes. So, Russ made a comment about popularity amongst infrastructure operators going back quite a while. So, I was saying that that seems like a useful thing to maybe include in the landscape description of Section 2. I would be interested if you happen to know or if we can find out what goes into this preference because that would be something for Section 3. Is it a characteristic of this software that it somehow appeals to operators, or has been appealing to operators? And if so, what is the characteristic that makes this so? And then specifically for the subsection that we're now describing, is this a characteristic that has—is it SSR related? Is there something security-wise that does so.

JOHN EMERY Okay. We did have two questions from Farzaneh in the chat, and I know that Peter just answered those as well. So, if you want to address those, Matthias.

MATTHIAS HUDOBNIK: No, actually I just thought could be one point or points, considered in the chapter related to incident response and patch management where you would say you have some comparison, commercial versions versus open source where you would say, okay, maybe of zero-day vulnerabilities, where you would say, whatever, a commercial software is maybe more rapidly or faster in the patch releases than for example, community-driven project.

MAARTEN AERTSEN: Yeah. Thanks for this comment. This is already part of the Section 4 SSR, which we didn't get to yet in the agenda. But I think it's a good comment. And I think we don't have text, but we do have the bullets. So, we had an SSAC member earlier describe that in their operations of a DNS service, they prefer open source because the patching speed is high for the software that they use. But more importantly if a vulnerability is in a dependency, they can actually check themselves if the patch has reached them yet. But if you have like a proprietary software, you may know that a component is vulnerable and is being patched not as transparent, and how fast this fix is reaching you.

You have to kind of rely on a change log or something, whereas with open source, this whole chain is transparent. And if you want to, you can actually put the patch in yourself on the day of whenever you have this knowledge. So, I guess, this kind of goes into—the instant response angle is new to me, so maybe this is something that we could discuss if we have sufficient input data. Thank you.

I don't read the chat myself but maybe there's—

JOHN EMERY

Yeah, you can if you want to come up and ask the question. Go ahead. Since you're here might as well.

FARZANEH BADIEI

Hi, I'm Farzaneh Badiei, non-commercial stakeholder group. I have been doing some research, which is very related to the open-source DNS, and I focused more on DNS resolvers. And I'm sorry, I don't know if in your document—unfortunately, I missed your meeting on DNS blocking. I had another meeting. But I'd be interested in knowing that because in the research that I undertook, it was very clear that different regulatory frameworks and court orders, they actually affected the open-source software like DNS resolvers that were operating on open source because they didn't have access to funding and they had to, all the time, comply with these court orders and regulatory frameworks.

I just wanted to know if you found something like some kind of regulatory framework or development that could affect the operation of the open-source DNS and could hamper it in some way.

MAARTEN AERTSEN

So, let me try to paraphrase to check if I understand the questions. So, are you saying that operators that use open-source resolver

software are less likely to have legal resources to respond to government requests? Is that kind of what you…?

FARZANEH BADIEI: In some cases, yes and in some cases, when they don't have—because open source doesn't necessarily mean free, right? Doesn't necessarily mean non-commercial. They can be commercial, but in some cases, when they are open source and not-for-profit like some regulatory framework and as you know, there is many. Day after day, we see copyright infringement blocking orders especially in Europe that is addressed to DNS resolvers. So, I was wondering if this could have an effect on maintaining open-source DNS product.

MAARTEN AERTSEN: So, I guess, this…

JOHN EMERY: Sorry, had a direct response, and then Maarten. Yeah.

UNIDENTIFIED FEMALE: Yes. So, I think from what I'm understanding from your question you're asking about regulatory frameworks in terms of—my understanding from the work I've done in Ofcom is there's some stuff around data protection. That if the regulatory framework has an issue around data protection, it's much harder for open-source

software generally. Yeah, so that's—is that your question? I'm not sure if I fully understand.

FARZANEH BADIEI: Kind of like, are there parallels. So, for example, we have regulations that could create struggles for people to maintain open-source DNS and to provide their service. Are there regulator frameworks that you have found out about that have had that impact?

MAARTEN AERTSEN: So, I think, it's not as common that the operator of the DNS service also is the developer of the DNS service. There are examples. For example, CZ.NIC operates to Czech ccTLD, and they operate not resolver and not authoritative name server. I always mess up that name. Anyway, so they do both, but I think for the other cases, it's not as common that the operator is also the developer.

So, when you tie together these two aspects of an operator being served some kind of court order or like regulatory request, I don't think it necessarily affects the developer. I'm also not super sure if the fact that an operator would use open-source software, if that translates in them somehow being more or less—so I guess, my answer is I'm not sure but these two things I think are not connected unless you are like CZ.NIC or like…

FARZANEH BADIEI So, shouldn't the open-source software developer provide some kind of capabilities to comply with these court orders? For example, have blocking and filtering in the—[inaudible] does that, for example.

MAARTEN AERTSEN So, I guess, that takes us a little bit into the DNS Filtering Work Party, but I guess if you have a resolver or an authoritative, you can just serve responses and some of these responses might be a like a blocking response or the zone may like omit some data or whatever, but maybe we should take this offline, if that's okay.

JOHN EMERY And we can coordinate to talk about the DNS Blocking Work Party as well. So, I had Matthias, Rod and then, Barry.

MATTHIAS HUDOBNIK No, just very quickly a response to her question. So, I think, from a privacy or GDPR perspective, I think it makes not a big difference or there's no difference if you're like a corporate or if you're whatever, organization as long as you are falling under the legal application of the law, you need to stick to it. And then of course, maybe, whatever, not-for-profit organization has maybe less capabilities due to a smaller budget or whatever, but I think this does not mean that the law would not apply on the same way if it's falling into the application of it.

ROD RASMUSSEN

Yeah, and I think, this is an interesting concept, where if you think about it, this is a product requirement in a way, right? If you think about the IETF, etc. DNS specs or this, that means the resolvers need to do these things as functionality. We have other market imposed, in this case, this is a regulatory imposed feature that would it requires development, right, in addition to the software. Now the question for the open-source developer is, do I want to provide that feature set or not. And market demand may be the thing that drives that. If you've got a lot of people saying you need to do filtering of some sort, then people are going to naturally include that in their product. It's not an IETF/technical thing, it's a policy thing, but it's a product feature requirement.

I think one of the good news on this particular issue is, you've got RPZ and kind of an open-source framework for doing exactly this kind of thing. So, there's already an open-source solution to this. And then there's questions about performance and things like that. This is way out of the scope of probably what we're talking about here.

JOHN EMERY

Barry and then Maarten.

MAARTEN AERTSEN

No, I'll just give a thumbs up. I think this is an interesting conversation, but I also think it might be a bit out of scope for what we're currently working on. But I'm happy to be convinced that it's not, and maybe we should take that offline if that's okay. Barry?

BARRY LEIBA Yeah, I want to go back to—what I heard you say a little earlier was that with open source, this was relating to knowing when a fix was coming and having control of that. That with open source you have transparency into who's doing the fix and when it's going to come, and when you're installing it, and you don't have that with proprietary software. That's what I heard you say. Is that correct?

MAARTEN AERTSEN So, I was paraphrasing input we received as a work party from an operator that is also an SSAC member.

BARRY LEIBA Okay.

MAARTEN AERTSEN So, this was one perspective that in their operations, they greatly preferred running open-source software because this enabled them to do that. Once their security team would learn of a vulnerability in a core component of their service, they would be in full control to get the fix rolled out in their service as opposed to if it would be a component they would even—the opposite situation would be where like their name server was proprietary, but they would maybe get from the vendor a bill of material style thing, where they would know what component is in there.

I suppose like OpenSSL is a component of this proprietary name server. Suppose OpenSSL lets the world know that there is a

vulnerability, then for the proprietary equivalent, they would not be in complete control of integrating that fix into their service, because they would need to wait for the vendor, and they would need to trust the vendor to disclose that update would have taken place. Well, I guess, you can look into the binary, but that was basically their point.

BARRY LEIBA Okay. It's the opposite of my experience with this. With open source, yes, if somebody makes a fix, you can see that. But nobody's responsible for doing it. There's no guarantee that anyone will do it. You may have to take that on yourself, so that there's that aspect, and that relates to what Geoff Houston was talking to us about.

MAARTEN AERTSEN Yes.

BARRY LEIBA And then the other side of it is that a properly supported software would have that information. I used to work for IBM, and IBM would tell you what fixes were in a particular set of fix batch or whatever, and they would disclose all of that and tell you how to install it and distribute the fixes. So, I don't know what kind of software whoever made this comment was talking about, but my experience with properly supported software is that that's better in that regard than open source, so.

MAARTEN AERTSEN So, this is why I would like Geoff's input on Section 3 because I think we should—

BARRY LEIBA Absolutely.

MAARTEN AERTSEN Yeah. This document needs to cover all sides, and right now, what I've noticed is that we got quite a lot of input that considers a particular perspective, which may be applicable to the major software systems in DNS, but might not be representative of all of open source. And I think this document needs to do both. It needs to point out positive and negative in general, and then maybe zoom into DNS, and clarify where some of the problems that are true in general may still exist in DNS.

And be clear about which problems do not appear to be a problem in our particular community because otherwise, it's not an honest perspective. But I am looking or we are looking for volunteers to plug in that realist perspective. And if you are interested in contributing, that would be very helpful. What I want to avoid is that we get this discussion started once we go near to SSAC review because that seems wildly unproductive.

BARRY LEIBA Sure. And I think it also speaks to probably a recommendation that we can make in here to software vendors that they expose this

information about security fixes, so you can see what security fixes are available, and actively choose to deploy them when you know you need those fixes.

MAARTEN AERTSEN: Right. So, to be clear, the comment that was made was about speed. It was not about whether your vendor was being honest or not.

BARRY LEIBA: Okay.

MAARTEN AERTSEN: Please.

JOHANNES LOXEN: Hi, I'm Johannes from SerNet. I work in the open-source regime for now more than 30 years. We work on Samba SMB implementation in Linux as a company too. And I want to remind that open source is not about the free beer, but about the free speech in that if you want to use a software professionally, you always have to take care that your software is running rightly. So that means if you want to buy from a vendor a closed source software, then you have the attitude that the software has to be right and it has to be fixed fast.

But if you want to have software for free, you still have to take care. And if you like the speed and everything, you ask the one who uses the software. You still have to fund your open-source software

providers to be able to get those fixes done as fast as possible. So, this is always a question, who funds which work and who is responsible. And if the open-source developers are not responsible for the work, but the user who uses that, that is difference from the customer vendor situation. And the last remark, SBOMs. We see now that open source has perhaps a bit better chance to get reliable SBOMs than some proprietary vendors because not all documentation in the proprietary world is as good as you think.

MAARTEN AERTSEN

Right. So, thank you for the comment from someone outside of SSAC. So, for those not aware, an SBOM is a software bill of materials, yeah. Are there any further comments or inputs from SSAC on 3.1 as a Section because otherwise, we'll move on to 4.1, which is exactly the same exercise, but it goes to the operational aspect. So, we just discussed characteristics of open source and security or stability properties that relate to those characteristics, and this section is meant to go about the operational perspective of using open source and implications.

So, for example, the whole patching speed discussion was a bullet in this section. So, I guess, we've not been super honest with the agenda here. But I'll give you a minute to think about this perspective. Oh, I'm being reminded by our wonderful staff that this work party meeting features opinions of individual SSAC members, and in some cases non-SSAC members that are invited to bring forward their opinions, but because this is a work party and

not a finished product, the things we're discussing are not SSAC consensus but how the sausage is made.

So, I guess, me personally but we as a group are still getting accustomed to our meetings being open during the ICANN event. But it also means quoting from meetings like this doesn't correspond to what we believe as a group. Because that's what you can read when we publish stuff. So yeah, thanks for pointing that out, Danielle.

And then, for this section, I'll look at one operator that I know runs open source. Peter, tell us about the implications for security and stability. Is there anything you like about using what you're using? Can we generalize your experience into this report? I'm sorry for putting you on the spot.

PETER THOMASSEN That's okay. This is Peter. That's a good question. Overall, I think, open source is more transparency in terms of what the code is obviously because it's open. So, it's easier to assess the code for vulnerabilities. Can't necessarily do that for proprietary software, and you have to hope that the vendor does it. On the other hand, obviously, it's also public. If there is some vulnerability to be found that has not been circulated yet, somebody finds it, doesn't talk about it, and exploits it. The view that we have at the SSAC, which is that operator, is that the benefits of open-source security outweigh the costs.

We not only use DNS specific open-source software, we use all kinds of—we have a Python backend application that uses various libraries and stuff. And often, you import a bunch of dependencies, and then random stuff happens. So, what we try to do and it's very hard to do this very thoroughly, is to actually read code diffs of different library releases before we deploy that for critical stuff. We do that. Otherwise, we change locks and hope for the best. For example, for the database library, the [Postgres] Python connector, I don't read the code and it's used by so many people that I kind of hope that other people would notice.

And I guess, I think we're already taking quite a proactive stance here. Most people probably don't read any code diffs, so that's how we handle it. I don't think it's a reasonable expectation though for open-source users to do that, so we shouldn't be recommending that. I don't know what you're taking away from this report.

MAARTEN AERTSEN Would you say there's some kind of trust relationship you build with the people that make the software that you use?

PETER THOMASSEN Not automatically. I think the trust relationship is a function of having interacted with them, and specifically having seen folks' competence or incompetence. We've also particularly not used certain libraries and written replacements for our platform ourselves because we found that key handling was insufficient or something.

MAARTEN AERTSEN Any other comments about SSR implications in operation of open-source software?

PETER THOMASSEN Maybe here's one more. So, if there's a security issue that gets discovered, either you discover it or somebody else, it is at least, in principle possible if you find that you're affected to also fix it yourself. With proprietary software you can't do that, so if you are willing to put in the hours or the money, you can actually protect yourself. I think that's a big difference.

MAARTEN AERTSEN Yeah, thank you. That's a good point. I just put in the bullet. I almost managed to spell the word yourself. Let's see. If there's no more comments, then I think we will move to the last point of the agenda and add time or maybe a little bit early because we are on a timeline. And what the idea is to finish most of the text of this document by the end of this month, and then do like a review phase internally. And if we deem that this review phase is of sufficient quality, we'll do a first look within the broader SSAC. But that's kind of a challenging timeline because as you saw Section 4 is not there yet.

The other sections are way better, like way further along. But that's what we're aiming for, and this whole timeline is assuming that we manage to publish by the next meeting, but that will be challenging. Well, depending on engagement. How many people

show up early enough to share their opinions. So, I really appreciate a number of you joining today who have maybe not been able to join all the calls. And I would welcome your contributions in the document.

Next work party meeting will feature a discussion of the survey we did. It was open for a month. There were about 100 respondents, which was more than we expected. And I think we got some nice input that we will be able to weave into the document. I guess, this is where I asked if there's any questions from the floor about this work or the way we're headed right now, and I'll first do a round of SSAC. Any questions? No, great.

ROD RASMUSSEN Just a comment, that's a super aggressive timeline speaking from a lot of experience at this. I would frankly even instead of—if I was in charge, I'd be saying, why don't you look at the fall meeting for this rather than this summer meeting, but if you can get it done, great.

BARRY LEIBA The answer for that is the other way around. We would rather have an aggressive schedule and miss it, than start off saying we're not going to get it done.

MAARTEN AERTSEN So, in the other perspective, and I appreciate what you're saying, because I know it's aggressive and it has been from when we

started is, I am interested in doing this work, which is why I'm doing it, but I'm not interested in doing this for years. So, if we don't get it out the door in a reasonable time frame, I will just admit that this is hard, and it didn't work this time. But so far, things have exceeded my expectations and I'm hoping we will continue to do so. But yeah, I'm not promising you any beer or anything.

ROD RASMUSSEN I was going to say, if you guys could pull this off by the summer, I'll buy you beers. That would be quite something.

MAARTEN AERTSEN Thank you.

BARRY LEIBA There's an incentive. You should take wine—wine from your wine region would be lovely.

MAARTEN AERTSEN So, if there's no other questions from SSAC members, are there questions or comments from non-SSAC members? And otherwise, seeing none, I'll mention two things that are coming up. So, first of all, John is working with ICANN to see if we can get a little bit more information about the—what it called? I'm sorry, jet lag is kind of kicking in. This is night time for me. About the…

JOHN EMERY It's killing me too, so one of the appendixes.

MAARTEN AERTSEN Yeah, yeah.

JOHN EMERY Data escrow.

MAARTEN AERTSEN Yeah, exactly. The report would benefit from having some statistics about data escrow use, and ICANN has that data because they receive reports. We're just trying to see if we can get something out of that. And then the other thing that is in the future is that, this weekend, the daylight-saving time switches here. And by the end of March, it switches in Europe. And from the 4th of April, which is our third work party meeting from now on, we will move to one hour earlier.

We switched for a while to this time slot at the request of Nabil and others, but from that time, I cannot make it myself, and that will be a problem. So, we're moving back. If there's a lot of complaints, we can discuss, but that's a little announcement at the end. Nabil?

NABIL BENAMAR Yeah, thank you. But now, it will work for me. It's okay, so.

MAARTEN AERTSEN	Oh, that’s great. We'll put this to the list too because there's some people not here today. And I'll thank you for your time, and I'll give you back about three and a half minutes.

[END OF TRANSCRIPTION]