ICANN85 Mumbai | CF – SSAC Work Session: DNS Abuse and AI Work Party
Sunday, March 8, 2026 – 16:30 to 17:30 IST

KATHY SCHNITT

Thank you. Hello and welcome to the SSAC DNS Abuse and AI Work Party meeting. My name is Kathy, and I am 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 today, this session is designed for the internal work and discussion of SSAC members. To join the speaking queue, SSAC members must be logged into Zoom and use the Raise Hand feature. When called upon, please state your name for the record. For observers, please note that the comments or questions in the Zoom chat will not be read aloud. However, the chair of the session may invite questions from the audience at their discretion and time permitting. And with that, I am happy to hand the floor over to Daniel Gluck.

DAN GLUCK

Thanks, Kathy. Welcome, everyone. This is our first work session at an ICANN meeting, so we are looking forward to that. We have a very simple agenda today. And by very simple, I mean there are two items: the charter and the outline. The charter was sent to the work party for their final review last week. There have been no

comments received via email or within the document this past week. So I will now give it over to Matthias to get us started.

MATTHIAS HUDOBNIK Thanks a lot. Hello, everyone. Just again as a reminder, it is a working session of the SSAC DNS and AI Work Party. We decided to have it public so people can participate and see. So we are at the start, but basically, it is just open to comments for members of the work party, just so that people are aware of it.

First of all, I would like to clean up the document. As Dan already said, we had the period until this Friday. There were no more comments. So I would like to go through and delete the points we made. And then the idea is to start with the outline and content work related to the definition of AI. I do not know. Are there any questions or comments from work party members? Perfect. I do not see any. So, okay. Then we can just start, exactly, to put out the points.

DAN GLUCK Since there were not any changes?

MATTHIAS HUDOBNIK No, there were none.

DAN GLUCK Do we want to accept the changes and read it without the brackets?

MATTHIAS HUDOBNIK Yeah, I mean, we can still read through the changes if you want. That is actually a good point. Let's do it. I would expect that people read it, but maybe let's start with the point with the...

So, we had some discussions related to the first sentences. What we changed it to: work on the intersection of AI security and the DNS is useful at this point in time because AI tools may be able to exploit weaknesses in the DNS ecosystem we set, and also giving an example such as domain registration at large scale and speed far beyond conventional human capabilities. Here again, we made the points at this point in time because the change will be rapid and also the relation to the DNS ecosystem and not just DNS. So we can delete... Is there... Billy?

BILLY Does weaknesses in the DNS ecosystem such as domain registration, phishing, and spam sound okay to you? Because it looks like we consider domain registration as a weakness in the DNS ecosystem. Does it not look weird?

DAN GLUCK I think we can accept the changes then reopen things. Are we reopening the charter?

MATTHIAS HUDOBNIK Yeah. No, but because we are not re-

DAN GLUCK We are trying to review it, but I think if we basically accept the changes, accept everything so we can read it clean. Got it.

MATTHIAS HUDOBNIK Yeah. That's exactly it. Because-

DAN GLUCK Billy's point is, I think, some of the terminology, once you take all the changes in, we may have to tweak them a bit to make sure it is still clear.

MATTHIAS HUDOBNIK Yeah, we can do that, but I do not want to open the charter discussion again because there was the review period twice. So I mean, we can do that that we-

KATHY SCHNITT Matthias? Can we please make sure that everyone speaks into the microphone? Otherwise, it does not pick up on the recording. Thank you.

MATTHIAS HUDOBNIK Okay.

DAN GLUCK Doing that.

MATTHIAS HUDOBNIK Perfect. So then back to your point, Billy. We had to change to: work on the intersection of AI security and the DNS is used at this point in time because AI tools may be able to exploit weaknesses in the DNS ecosystem, such as the domain registration. Here is also like there needs to be after registration the comma deleted. One up-

JEFF BEDSER So Matthias, I think that with the charter not being the final work product that goes out for publication, we can worry about some of the weakness in the language we have done here of interpretation. But I do not think that, to your point, it is locked. We know what we mean.

And we can start the document and start working on the next part because while we could keep tweaking this to make it clearer and clearer, this is never going to be published as the charter. The charter is for our internal purposes to say what we are working on.

MATTHIAS HUDOBNIK Yeah, exactly.

JEFF BEDSER So.

MATTHIAS HUDOBNIK Yeah, that is a good point. So then let's go again to the in-scope, where we start with the first point like establishing the preliminary definition of AI and scope of analysis for the SSAC. So here we have the AI definition outline. Can you... Rick first.

MATTHIAS HUDOBNIK Okay. Rick?

RICK WILHELM Thanks. Rick Wilhelm. Actually, this is just about the deliverables section a little bit up, and this is just a minor point that probably I could have brought up in a comment somewhere. But in the first paragraph under the section heading Deliverables and in the first sentence here, "Intends to produce multiple work products," then we say, "Deliverable one," and then later down we say, "Potential future deliverables."

I do not know that the way the first sentence is written, where it says, if you scroll up just a little bit, "Intends to produce multiple work products," seems contradictory with the later sentence where it says, "We may produce additional things." And I do not know if... I would say that we should not have the first sentence in this section firmly commit the work party to producing multiple deliverables, but instead... because actually it is pretty rare for a work party, in my experience, to produce multiple deliverables. That is pretty rare.

JEFF BEDSER Do you like that fix? Do you like that fix, Rick? The blue word I just inserted.

RICK WILHELM "Intends to potentially produce." Sure. That is probably enough equivocation, so thanks.

MATTHIAS HUDOBNIK So thanks for the point, and maybe just to react. Actually, we had quite some discussion related to the charter back at the beginning. So the idea was to make three deliverables, and now we said, "Okay, we will just focus now on one," because we reiterated the whole charter, and that is why this is still there so that people are still aware. I do not know if you were in this session, but at least this is the idea so that we have now, as we say here, a short educational one and then there might be, or most likely we will have one or two more. And actually the plan was to have a bigger document where we would have all of these three discussed. Okay. Are there any...

NABIL BENAMAR Matthias, I just want to add a clarification for Rick. So when starting talking about this and debating, we have seen that AI is being used for attackers and for defenders too. So surveying on the AI landscape might tackle only the attacks, maybe. The AI-generated attacks, different speed, different scales, et cetera. And then we thought that maybe this will be too long if we include two, so this is why we might go into two documents. So maybe it is rare.

Definitely, I agree with you, it is rare for a work party to, but we were afraid that it would be too long to include everything.

RICK WILHELM

Oh, I agree. Yeah. Sorry, into the mic. Two fingers? Yep. Rick Wilhelm again. I agree, and I recall the discussion that we were very worried about the document getting too long, and that continues to be a concern. And I just wanted to point out, and Jeff made the edit, I think, to just equivocate a bit in the document. That is all. Because I am agreeing with all that stuff.

MATTHIAS HUDOBNIK

Thanks, Rick. No, I mean, that is totally fine, and I think it is also good. The more brains work on it, the better it will get. But again, maybe just to summarize it once more. So the idea is really, as we said, we have this short educational reference report where we define what SSAC thinks about AI. Then we have the categorization as you see in the scope. And the last thing would be then analyze the current most noticeable or likely impact and also projected security threats. So and that's basically the idea. And then as we said, there might be further deliverables based on the need or based on the ideas we see or we come up with out of this working party. Are there any other comments? Perfect. If not, so then can you open the-

DAN GLUCK

The outline-

MATTHIAS HUDOBNIK

-the outline document? So the idea now is to start the content discussion related to establishment of a preliminary definition of AI and scope of analysis for the SSAC, and also clearly identifying which technologies are within the remit of SSAC's work versus AI more in general. And we should also keep in mind when we start with this discussion that we have then also some kind of catalog examples as a further point.

So the idea or my understanding would be that we, as we already have... when you go a bit down, we have some definitions already from the ISO standard. There are two definitions, and then we have the EU AI Act. And here my understanding would be to really take one of these definitions which are already established, and then we could tweak it as we need it for our SSAC remit. And then related to the categorization, there are also some kind of general categorizations where you say, okay, you have on the one hand artificial narrow intelligence, you have artificial general intelligence, and you have artificial superintelligence. We will just focus obviously more on the narrow intelligence, which is like machine learning, neural networks, deep learning, and then also GenAI, agentic AI includes also these areas.

And so the first thing what I want to throw into the group is the ISO definition of artificial intelligence systems also like the AI Act, and it is defined here as an engineered system that generates output such as content, forecasts, recommendations, or decisions for a

given set of human-defined objectives. So this is in this standard. And then we have... Can you go a bit down then? I think there should be also the definition of the AI Act.

RICK WILHELM

Rick. Can you go back north? I am looking at the first bullet, and I am wondering, does it matter for the purposes of the rest of our subject matter does it matter what the definition of AI is for the purposes of what we are getting into? And I am not just asking you, Matthias, but asking for this paper, does it really matter? So what? As an old boss of mine used to say. Because when we get into the second bullet, legitimate and illegitimate uses of AI in the context of DNS-

MATTHIAS HUDOBNIK

So sorry just to interrupt you. Just forget the legitimate and illegitimate use of AI. It is like this because here the work plan and the charter are things we were putting out of the charter before. The idea would be really when you go back to the scope document, the catalogs... So can you just quickly go back to the other document? Actually, these three bullet points are the things we want to emphasize in this document.

Now the first section is, as I said, establishing a preliminary definition of AI and scope of analysis for the SSAC. So that we say we take a definition, and then the next point would be in the document to catalog examples of how AI is used in both offensive

and defensive manners. And then the third point would be the analysis of the current most noticeable threats. And I think it makes a lot of sense to define AI for SSAC's purposes because a broader definition in, for example, the EU or the ISO standard might be something different than what we are thinking and what we think is relevant. We can also come to the conclusion that we are not defining it, but the idea would be of this document to do it because it is an education document, and we can refer also to it then and say, "This is what our understanding would be."

RICK WILHELM

Right. So to the extent—I am just saying that does it matter for the sake of what we are... academically it would matter, but for the sake of what we are trying to accomplish in the context of DNS abuse, does it matter? And I understand what the words that are on the screen right now say: establish a preliminary definition. But that is just there because we wrote it there. That could be taken out if Jeff Bedser twitches and deletes it just like he made that edit. So that is not canon either. It is just there because. So I am just asking. And the only reason I say that is because I am trying to make the task that is in front of us doable and easier and not trying to make it hard. Because it would be a moving target, and we would almost never finish because the definition of AI is moving so fast. So let's just solve a smaller problem, as another boss of mine used to say. That's all. So now Warren's got his hand up.

WARREN KUMARI

Yeah, I mean, people are not going to be surprised I am agreeing with Rick. I think that we need to get this document done and out the door as soon as possible, and we could easily spend five or six meetings talking about the definition of AI. Right, because there are literally thousands of different definitions, and no matter which one we choose, somebody is going to want a different one.

JEFF BEDSER

So, I think in one of the earlier calls, the point was made, to your point, Warren, that there are multiple definitions. But for a baseline anchor, much like we had to do with DNS abuse, we need to basically say, "This is the definition we are pulling from, and everything we are talking about is based on this. Yes, we understand there are other ones." So I do not think we need to define it. I think we just choose a definition that has already been made and that fits our purposes and use that as the baseline, so someone does not question our paper by saying, "Well, you could have meant this if you use this definition and this and this." It is anchoring it to a particular definition. So hopefully, we are not planning on creating our own definition. Hopefully, we can just choose one and use it.

MATTHIAS HUDOBNIK

Yeah. I think that for any report that includes some overview about some technology, we need to agree on a definition of DNS abuse and also of AI. So that was the intent of this. Because AI has different definitions, so we just need to agree on one to put in this

document. There will be an overview section in the top before we dive deep into the details. But I agree with you, it will not make any difference to adopt this one or the other.

WARREN KUMARI Peter has his hand up.

PETER THOMASSEN Yes. So the three bullets that are in scope mention attack or defend the DNS or the identifier system, and they mention impacts and security threats on the DNS naming identifier systems. They do not mention DNS abuse. DNS abuse is not an attack on the DNS or the identifier system, I would say.

But the work party is called DNS Abuse and AI. So do we deal with AI and how it affects whatever attackers or malicious parties' efficiency in conducting DNS abuse, which typically is understood as phishing, malware, spam, and these things, or do we deal with how AI offensive and defensive tooling might or might not affect the DNS naming identifier systems? I think we need to align the work party purpose with what is in scope.

MATTHIAS HUDOBNIK Thanks, Peter. And next, Warren.

WARREN KUMARI

I took myself out of the queue, but what I was going to say is we could just say AI without defining it. Readers would know what we meant, but it is not worth having the discussion.

MATTHIAS HUDOBNIK

Okay, yeah.

JEFF BEDSER

I think Peter made a really good point. The way we have worked out this scope, it is not addressing the way DNS abuse is being hit by AI by saying about attacks to the DNS and identifier system. So if we could modify the name and just say it is AI and the DNS system and attacks in the DNS system, or we can tweak this a bit to make sure we are also covering the use of AI to facilitate attacks via DNS abuse. But Peter has identified a pretty big hole between the title of our work party and the way we have defined the scope.

MATTHIAS HUDOBNIK

Actually, the idea was to tackle both. And again, to the definition question, as I said before, the idea is that SSAC has some kind of, as we said in the Deliverable One, short education reference report establishing the context, definition, and use cases. And here again, to start, we should have some kind of definition of what we think about AI because it is important; otherwise we are not talking from the same angle.

And again, this does not need to be very long, and we should not define it new. So we can use an ISO standard, can say, "Look, this is

what we think about." Maybe we use less of this definition because this is not relevant for the SSAC's purpose. But I think it is important because otherwise, if you say, "Okay, I do not care," then I can ask you, "Okay, Rick, what is AI for you?" And you might say, "Yeah, for me, it is machine learning." Another person would say, "Yeah, for me, it is a recommender system." Another one would say, "Yeah, for me, it is agentic or the use of AI assistance." So we have a lot of confusion then.

RICK WILHELM

But I think that—and I want to give a shout-out to Andrew because I think he phrased this well. I mean, I just said it does not matter and sort of shrugged my shoulders. Andrew put into the chat as an observer: the definition of categories of AI does not matter to attackers. They select whatever tools serve their purpose at any given moment, right? And I think that... look, it literally does not matter. And us defining AI in our definition whether or not we reference an ISO definition—which if they have made a definition, I will bet my lunch money tomorrow that it is out of date.

Whether or not we define what type of AI, whether it is super intelligence or Superman intelligence, or Batman intelligence, or artificial general intelligence—it does not matter for the purposes of this paper. I am just saying that by us defining it, it just gives people a potential avenue to criticize the paper. And so, let's just say that we are talking about AI for whatever the current definition is. But if you don't want to take that input, I don't care. That is fine.

So that is the point. I am just trying to make the paper easier to write, easier to edit, and quicker to get done, because I am trying to make the paper smaller. That is it. Thank you.

MATTHIAS HUDOBNIK Okay. Thanks, Rick. Now to Georgia, and then-

JEFF BEDSER Can I just do a quick response? So, I think to move this forward, the simple solution here is if the paper as it produces needs a definition, we can always put it in later. There is no point in actually spending too much time when half of us believe it should have one and half of us believe it did not. Let's put that aside.

Later on when the paper is close to completion and you are reading it and you are like, "Oh, crap, we really do need to define it," we can put it in then, and also we will be six months from now, and maybe things will change enough that that definition will be valid. But to the points about it keeps changing, well, if we choose a definition right before we publish, we are more likely to be right if we need one. So I suggest we put that discussion aside and move forward.

But I think we need to get back to Peter's point, which is "in scope" needs to be something about DNS abuse, where the AI is being used to facilitate criminals using AI to facilitate volumes of DNS, which is part of the charter. That is not in the scope now, and it needs to be in the scope. So we need to add a bullet unless someone disagrees, because otherwise the work party we agreed to join is not covering

the topic we were planning on covering. So I think we should add one bullet referring to criminal use of AI to facilitate at-scale DNS abuse.

MATTHIAS HUDOBNIK: So sorry, now we open another discussion. I think I would like to stick to the definition one, and then we talk about this criminal use because we have here offensive and defensive manners to attack or defend the DNS and identifier system. And I think this-

JEFF BEDSER: Right, but DNS abuse is not attacking the system. DNS abuse is about attacking the end users. The problem is DNS abuse itself is a poorly termed definition because DNS abuse does not abuse the DNS. So we do not want to continue using that. The DNS abuse we are talking about is criminals using domains to facilitate phishing and malware distribution, which is what DNS abuse is about. It is not an attack on the system. The system is completely intact no matter how much DNS abuse happens. So that is why we have to define it separately.

MATTHIAS HUDOBNIK: Yeah, thanks for the point. Now to Georgia.

GEORGIA OSBORN: Yeah. Mindful of moving the conversation forward, I think putting on my day job hat, it is important to have a definition that can stay

consistent throughout the whole of this, which is why I am going to caution against putting a definition there, especially for future documents that we might want to look at in the future. I think no matter what you choose to put there, a definition or not, just be mindful of anything that we might add to the future. So perhaps better to not put one at all.

NABIL BENAMAR

Just a clear point about when we started working on this. We have established that we want to say to the readership, in the case of DNS abuse or DNS security to make it general, what has changed with AI? That is the main point. Because DNS security is there, DNS abuse is there, we are not really defining this, but we want to see what has been changed with AI in terms of speed, velocity, and other stuff. That is it.

MATTHIAS HUDOBNIK

Any other points? So I just want to add two things. First related to the definition. I need to kindly disagree with you, Rick, because the ISO definition is there since more than five years, and it is not changing every day or quite quickly. For me, it is more that we talk about the same thing because when we have discussions in the community, I think it is helpful, especially as we said, this should be a first educational document where we defined some definitions, some scope, and also the use of AI in speed and scale, and also to abuse the DNS with it, but also to use it for legitimate purposes.

And that is why we had also in the second bullet point some categorizations of examples on how AI is used in both offensive and defensive manners to attack or defend the DNS and the identifier system. And we had also at the third bullet point the last time some examples related to registration data, which we put out because we agreed on that and said we do not want to be too specific because we want to see how we further evolve in the document and that we get it out faster.

My opinion is the definition would not change or would not make the document much more complicated because at the end, it is just that we talk about the same thing. And it is still possible to change it at some point. But as we see in most standards, there is not a big change. It is just that you know what you are talking about. Because in the discussion we might have then, we say there is an attack driven by AI or by AI system, and then somebody might say, "Yeah, but what do you think about this? When we talk about an AI system, what is it? Is it machine learning? What is it?" And that is why I think it is important to have some kind of reference. Nabil?

NABIL BENAMAR I remember that recently in Davos at the World Economic Forum, one speaker said that we need to stop defining AI as a tool. It is an agent. And this is a big difference between the two, which means that an agent is independent, can think, can take actions independently without giving back to humans. And this changes drastically all the landscape. But this is out of scope of the current

document. I think we can move forward, and finally, when we need a definition, we can get back and add it.

MATTHIAS HUDOBNIK Thanks. Any other thoughts? Okay then.

DAN GLUCK Hold one second. This is Dan from support staff. I just want to understand from a secretariat staff perspective, reopening the charter, if we were to go through these two things-

MATTHIAS HUDOBNIK Yes.

DAN GLUCK -that will add another week of full SSAC review.

MATTHIAS HUDOBNIK No. So the idea—there was now the deadline, and people could have made their comments. So I would be very much in favor to stick now to the charter because we changed it several times. And also, I would encourage people that they really attend the sessions, not come in once or twice in a month or two, and then we should change it again.

So I think there were now two times the possibility to reopen it, or we reopened it, and I think we should try to start and move on with

the document, and we can still change some things in one way or another. But this would be my understanding.

JEFF BEDSER

So Matthias, based on Peter's point that I emphasized, in the scope is not the top—nothing in the scope covers the topic of DNS abuse in the way that DNS is abused within the industry. We are basically saying that there is a hole, that we have not actually addressed the topic we are supposed to be discussing at all in that scope. And I think that we have to at least open a charter and get that back in because most of the people who joined the work party were joining it because it was addressing DNS abuse.

So I think that that is the one thing we do need to change, and if we do not change it, I understand it might add another week, but it seems kind of silly to have a party about DNS abuse and AI and not talk about DNS abuse in the work party.

MATTHIAS HUDOBNIK

Yeah. Noted. Then I would be very interested what would you or Peter suggest to change? Because we put out some wordings related to abuse, and we agreed on that. There were now twice the possibility to comment. Peter did it once. And then I also would like to give the floor to Gabe.

GABE ANDREWS

Yeah. Just as an example of what Jeff is getting at, but I want to make sure that I am not misunderstanding your point of view here

too. I have been collecting anecdotal case studies of where law enforcement is seeing criminals, we think using AI and new automation to rapidly increase the scale and scope of phishing campaigns and trying to pull in real-world evidence that I might be able to just share our perspective of what we are seeing. Would that be included in the scope as you have drafted, do you believe?

MATTHIAS HUDOBNIK I mean, if I remember correctly, we agreed this would be definitely... So we have the last bullet point when you go a bit up, which is now quite broad. So: analyze the current state, most noticeable or likely impacts or projected security threats of these technologies on the DNS naming and identifier system. And we had before some concrete examples related to registration data, for example. And for me, it is definitely possible to put here these kind of examples you mentioned. So I do not see any problem, and it is also not out of scope of the third bullet point. I do not know what... Jeff, what do you think? Please respond.

JEFF BEDSER Again, I think we had plenty of discussions in the charter formation that we were going to cover that. What I am basically saying is there is an omission in the scope, what's in scope and out of scope of that exact topic as just laid out by Gabe and myself. So I am just saying, my solution is a simple bullet. As I mentioned earlier, bullet number four would be the criminal use of AI to facilitate large impact DNS

abuses such as phishing and malware, and just add that bullet, and we are good. But I think Jothan has his hand up.

MATTHIAS HUDOBNIK Yeah. Jothan.

DAN GLUCK Yeah. Thank you. Everything that I heard Jeff say and everything that I heard Gabe say as far as some of the applications of DNS abuse that are occurring still fall under the umbrella blanket definition that we are using for DNS abuse. You know, the four plus the one related to the three...

JOTHAN FRAKES For other ones. But we have got something that is already defined as DNS abuse, and we could spend an awful lot of time trying to relitigate that. But really, where we are focused is AI and how it can affect those things. And so when we say DNS abuse, what I hear, what happens between Jothan's ears, is we are talking about the DNS abuse as defined, which covers phishing, and it covers some of the other things.

And yes, the agentic stuff is going to make it perhaps more powerful or more believable or more credible to trick someone into taking that action. But it still falls under those same definitions. And so we could spend an awful lot of time on that, but I do not know that we need to. Thank you.

MATTHIAS HUDOBNIK

Thanks. So again, like we had before, the question related to criminal use and legitimate and not legitimate use, and then we had the discussion about the things that we said we want to keep the definition as neutral as possible. And that is why we also said at bullet point two the example on how AI is used in both offensive and also defensive manners to attack or defend the DNS and identifier system, which would exactly cover these cases.

Because what is legitimate and not depends on the local law, depends on the region, and also on the governance. And that is why my understanding would be that we definitely cover this. And when we also go up in the explanation on the second paragraph, we also say: work on the intersection of AI security and the DNS. Yeah, thanks. Go a bit up, please.

So what we also said: such as domain registration at large scale and speed far beyond conventional human capability. So obviously, this would also include in the assessment the abuse aspect. This would be at least my understanding.

JEFF BEDSER

Yeah. I think where we are missing each other, Matthias, is that even if DNS abuse as defined by this community and SSAC, which is the use of technical DNS abuse, phishing, malware, et cetera, were to increase by seven thousand percent, it is not going to stop

the DNS from working. It is not an attack on the DNS. It is not going to affect the security of the DNS.

This is about harms to the end users, which is why your definition, as stated, that says about the security of the DNS, does not cover DNS abuse, and that's why we want to add it. And that is basically what I think the group is saying. I think we just need to add that in because the DNS itself will keep on running no matter how much phishing happens.

MATTHIAS HUDOBNIK Yeah. Noted. Rick was next.

RICK WILHELM So I think that in section three of the charter, the words DNS abuse do not appear, but it does appear in paragraph one of the description of the deliverable. So if you go... I am not sure who is running the screen. Dan, are you running the screen?

DAN GLUCK Yes.

RICK WILHELM If you go north in the thing to where it says deliverable one, and right there what you have got highlighted, someone has the word abuse highlighted. It says, "With a focus on DNS abuse there." Oh,

that must be me. I think that is me that has got that highlight right here.

So I think it is covered there, right? I mean, it is unfortunate that the word DNS abuse does not appear down in the scope. It feels like an editorial omission. I do not think that it is a semantic omission. It feels like an editorial omission to me. I do not necessarily think that it is a big deal to...

I think that Peter's... I would almost put it as a friendly amendment. I think his edit there is pretty targeted in that bullet that he put there. Now, probably still would require the charter to go out for re-review. It probably resets the clock on re-review, to Dan's point. But that's just process. I do not think it is worth breaking SSAC process for that, but it does reset the clock on it. But I do not think it really changes the semantics. I just think it changes the syntax because it is kind of an editorial change. But I think that Peter's targeted edit is pretty decent. Thanks.

MATTHIAS HUDOBNIK Thanks, Rick. Any other comments?

DAN GLUCK Maarten, did you lower your hand? Yes, I did. Cool.

MATTHIAS HUDOBNIK So when we look at Peter's little add here, this would be like... I mean, as I said at the beginning, the abuse part is in the

explanation. I mean, obviously, if we put here the use and abuse, it is still clearer. But if we would change it now, we would have another week to wait. So I mean, yeah. Are there any other comments?

Okay then, related to the definition again. So the idea was to come back to the outline document and discuss this actually. And yeah, we also agreed to have some preliminary definition of AI. So I would like to still come back to the output document and still have a look on that and would be interested what the people think about it and if we can find some minimum consensus where we'd say, "Okay, this is something which makes sense," and where we can also refer to it and even might change it.

But the idea was really to have maybe one or two short paragraphs and that is it, and then we refer to that, and it should not take too long, actually. And this was always the plan when we started with the charter. Can you...?

WARREN KUMARI Yeah.

MATTHIAS HUDOBNIK Perfect. So then coming back to the ISO definitions, if you see this here, this would be something which is a definition valid since 2022, as you see here, and this would be something where we could refer to it and say, "Okay, look, this is what we will..."

Let's say if we talk about an attack on the DNS, or let's say abusing the DNS with an artificial intelligence system, then somebody would say, "Okay, this is what we understand under artificial intelligence or artificial intelligence system." What are the thoughts on that? And we could even say, "Okay, we take whatever parts of it and that is it," and then go further. But we have at least something where we could say, "Okay, we talk about the same thing." Warren?

WARREN KUMARI

So, I put it in the chat, but I guess if we are talking about the definition, I will talk about it. A number of people have issues with this definition. For example, if I say, "I would like to know if I should buy a boat," and I shake a magic eight ball and it says, "You should buy a boat," that falls within this definition.

It is an engineered system that generates output such as a recommendation, "Yes, no, I should buy a boat," for a set of human-defined objectives. Give me a recommendation if I should buy a boat, right? Like, this definition is clearly not fit for purpose for this use.

There are many other definitions of AI which are also not fit for purpose for what we want to do, I believe, but this one clearly does not work in this case, right? Like a magic eight ball is clearly not an AI, but it fits in this definition. So we could spend a lot of time choosing from all the different definitions and finding one that does fit what we want, like large language model or similar, but I do

not think that is going to be a great use of time. But this one does not work.

MATTHIAS HUDOBNIK Thanks, Warren. Any other thoughts on that? Rick?

WARREN KUMARI Yeah. I mean, other thought is this is going to be a lot, a huge time sink that is not going to get us anywhere, but I have made that point already.

RICK WILHELM Yeah. Chuck the definition part, as I have said before, and just save a bunch of time because if someone does not know what... If someone is reading this paper and they... because the definition does not matter, as I said before.

MATTHIAS HUDOBNIK Any other? So yeah, for me, it definitely my understanding would be it matters. And yeah, we can agree and say, "Okay, this definition is not good," and there are not so many definitions. There is one with it like in the EU AI Act. Then there is one related to NIST, and we could even take some more technical ones like in OWASP or so, where they refer to it.

And in most of the documents, or not, I think even in all of them, you need to define when you talk about what, because otherwise it

is not clear. And I mean, so what would be then your definition of AI, Rick? I would be very interested.

RICK WILHELM

It does not matter what my definition of AI is for this paper. That's my point, is that for the paper that we are doing about DNS abuse and AI, the definition of AI does not matter. It does not matter at all what the definition of AI is. It could be a magic eight ball, it could be ChatGPT, it could be anywhere in between.

For the purposes of this paper, it only matters that there is some technology which is probably better than a magic eight ball and something north on the technology sophistication scale of a magic eight ball, and probably more than I can gen up using code over a weekend or even by myself, and it is what people are using to do stuff, and people are calling it AI.

And so it really does not matter. It almost has to just pass the Justice Potter test, which is: I will know it when I see it. And the definition is by definition moving, because that's why they call them frontier models. And it used to be—and we kind of all know this—that what now is disparagingly known as "that is only ML" is people used to almost call AI.

And ML used to be the coolest thing on the block. So it really does not matter. And so what we are only thinking about is cutting-edge technology that people are using for these sorts of matters, and

that is all we have got to write about because we are not splitting hairs about what is AI and what is not.

MATTHIAS HUDOBNIK Yeah. Thanks, Rick. Billy?

BILLY In five years, if you read the article, you would be interested in what we meant when we said AI.

RICK WILHELM So in that case, then we write it saying, "This is written at a point in time," saying that, "The year is 2026, and we are writing with the current state-of-the-art." It is such an easy thing to do that even an idiot like me could write that text, right? Even ChatGPT could write that text.

So that is an easy thing to do to write it with: this is written against the current state of the art at this point in time, and we are not going to try to define it. Because there is a... It is almost a Heisenbergan or a quantum problem that the harder you work to define it, the harder it is then to nail down and define.

And so just try and leave Schrödinger's cat in his box and say, "It is in there somewhere," and you do not have to measure it exactly. Just say, "Look, it is AI, it is 2026, and this is what it was called then," and the reader will be able to figure it out.

And that is all we have got to do for the purposes of the paper, and let's just get on to the rest of the paper, which is going to be hard enough for us to do. That is not easy pickings in all that stuff. Thank you.

MATTHIAS HUDOBNIK Thanks, Rick. Maarten?

MAARTEN AERTSEN Yeah, I have a process suggestion for the group and specifically to you as chair. I understand that the definition has made the charter. You have given people time to react, and it is in now. I also hear that people really do not want the definition.

So my process proposal for you to consider is to move the order in which you discuss the work items that were agreed to be in scope and start with the second one and table the first bullet for now, so as to give that a bit more time to simmer and make progress.

And I think that will also allow you to eat that week of charter review and just get going. So that's my process suggestion, and I hope it is useful.

MATTHIAS HUDOBNIK Thanks a lot, Maarten. Any other points from others? So in... Yes, Robert.

ROBERT GUERRA

Yeah. So this is Robert. Maarten, I think that's a great proposal. I think just given the time and discussion that we are spending on it, changing those sequence. But I would still add a comment that was mentioned earlier by Rick, that when we do get to the AI part afterwards, we do specify the moment in time.

We should have some text there, and there should be a moment-in-time definition because this document will be posted online, and if someone reads it in the pace of the development and the changes in the terms, it is very important to flag the time-sensitive definition items.

But also for citing examples, we should specify the time and the context because this is an incredibly rapidly changing field, and so we should also mention that in the opening as well, just to help us because if it takes us time to write this, by the time it comes out, we then might want to go through review to go any of the time-sensitive elements, do we need to tweak them or not? But in short, I support Maarten's proposal to moving the order.

MATTHIAS HUDOBNIK

Thanks a lot, Robert. So yeah, the idea would be to adjourn our meeting because we have six minutes left, and I think we should take the discussion offline. So I will definitely consider the points you made, and then let's see, and we might restructure it.

I think that is the best way, even though we had quite some discussion about it. Also for our participants, public participants,

yeah, you can see now this working party is quite... there are quite some heated debates, but that is part of it.

And the idea is also to have this public so that they really see it is hard work, and we try to find consensus, which is not always easy, but we are always able to find them at some point, some consensus. And therefore, I would like to close the session now. Thanks a lot for this fruitful discussion and give the mic over to Dan.

DAN GLUCK Thanks, Matthias. I want to thank everyone for joining for the work party. Our next scheduled call is going to be Thursday, the 26th of March at 17:30 UTC. So, thanks everyone for joining for this window into the SSAC process, and you can stop the recording.

[END OF TRANSCRIPTION]