ICANN83 | PF – Joint Session: ALAC and SSAC
Tuesday, June 10, 2025 – 09:00 to 10:15 CEST

YESIM SAGLAM

Hello and welcome to the ALAC and SSAC joint session. My name is Yesim Saglam, and I'm the remote participation manager for this session.

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. During this session, questions or comments will only be read aloud if submitted within the Q&A pod.

Interpretation for this session will include English, French, and Spanish. If you would like to speak during this session, please raise your hand in Zoom, and when called upon, virtual participants will be given permission to unmute in Zoom. On-site participants will use a physical microphone to speak. Please state your name for the record, the language you will speak, if speaking a language other than English, and speak at a reasonable pace.

Thank you all for joining. And I will now hand the floor over to Jonathan Zuck, ALAC chair, and Ram Mohan, SSAC chair. Thank you.

JONATHAN ZUCK

Thanks, Yesim, and thanks everyone for being here. We always have a great collaboration with the SSAC. And so we always look forward to finding out what you guys are up to, what you've been researching, and then hoping to find ways to boil it down so that, well, first of all, that we understand it, and then secondly, that we can help our constituencies to understand it as well and help get some of that material out. So, thanks and welcome.

RAM MOHAN

It's always a pleasure to be with ALAC. Back in the day, it used to be a group where you had to spend time explaining the technology. Now you don't have to explain the technology. Now it's really much more about figuring out how to explain it to the folks, to the other folks, because there's a great deal of expertise here. We can do all of the jargon that we normally do, and we don't have to really explain it too much to all of you. But it's always a pleasure to be here.

We particularly look at the collaboration, Jonathan, with the ALAC as a crucial part of what the SSAC is doing. And a big part of it is because of a realization inside the SSAC that it's a very important thing for us to do the work and to provide research and advice and things like that. But the real value is after it's done, how do you get it in a way that is understandable, digestible, and in the hands of those who probably need it the most, right? So we're evolving from a place where we patted ourselves on the back and we said we've

done a good job writing a good report to saying that's when the job begins [audio glitch 00:03:30 - 00:05:27]

JONATHAN ZUCK -- word out and we talked about the phishing course that we're doing. But that same mechanism is one that should be ongoing. And I think, you mentioned the ALAC, but part of what's really unique about our community is that there's a greater At-Large community that includes these regional, At-Large organizations, these At-Large structures and individual members that create a kind of coverage around the world that I think is unmatched within the ICANN community. And so we're going to be doing lots of experiments to really understand how to leverage our largesse, so to speak, and doing that in partnership with you would be a great thing to do, for sure.

RAM MOHAN That's fantastic. Yeah, I'm aware of the ecosystem beyond the ALAC itself, and that's a real draw for us, because that gets things moving. And we have lots of ideas also, Jonathan, on ways to make the material that we are coming up with, make it more accessible, right?

So one of the things that we're doing is for many of the seminal SSAC documents, we're having our members come up with a 500-word executive summary that kind of captures the essence of what that is. And what we'd like to do, for instance, is to work alongside

you and figure out how we take that, maybe literally animate it, find ways to make it presentable, find ways to make it accessible to the audiences.

JONATHAN ZUCK I think Barry mentioned something about an interpretive dance or something like that, so we're looking forward to that, so thank you.

RAM MOHAN So the first topic is, the SSAC, just a few I guess weeks ago, we released a document, SSAC127. This is on DNS blocking revisited. And we have one of the co-chairs of that work party, Warren, here to present to you the key findings and some of our thoughts on that. This is an area that received a great deal of attention inside the SSAC, and we're engaging with you. We're also going to be speaking to the GAC and other folks. Because this is a, as you all know, we find policymakers, governments everywhere increasingly resorting to DNS blocking as a reflexive mechanism for various reasons, and we're here to both shed some light on it, but also provide a bit of caution on it. So, without further ado, Warren.

WARREN KUMARI I guess next slide. Who's doing the slide events? Anyone doing the slide events? Okay. Can we have the next slide? I'll just wave my hat at whoever's. Well, the next slide is really pretty. Has a nice

picture of Prague. We actually have a lot of slides to get through. So I will go relatively quickly.

RAM MOHAN That's it. Warren.

WARREN KUMARI It's okay, I can just talk in the meantime.

RAM MOHAN Warren, talk to it and let them figure that out.

WARREN KUMARI I'll figure it out. So basically, what is this? Back in 2011 and 2012, SSAC published SAC50/SAC56, which spoke about DNS blocking. Obviously, the internet has changed a whole bunch in the last 12 or 13 years. And so we decided it was time to update the document and provide some new advice, especially around things like the effectiveness of DNS blocking, etc. So, as I said, the slides are pretty. Next, next. Yep, there we go.

So, the purpose of the document is to primarily advise policymakers and government officials on how effective DNS blocking is, but also the dangers around it, sort of the collateral damage that can happen if it's not deployed well, etc. And next slide. So what is this DNS blocking that we keep talking about? Basically, it is a way to block access to content on the internet by

performing some sort of filtering at DNS recursive resolvers. The goal of this is to present a set of users from accessing that content. But what's important to know about this is it only affects users who are using the particular recursive resolvers where the DNS blocking is applied.

And what it's supposed to accomplish is, next slide, basically this. You shall not pass. So, if you are using one of the recursive resolvers where the DNS blocking is implemented, you are unable to access the content that is blocked. But, as noted earlier, this is only partially effective and only applies where the blocking is done.

Next. DNS blocking is just a tool. As with any tools, it has some strengths and it has some weaknesses. And understanding how the tool works is important. Because if you don't, you're liable to cut yourself with it. It is potentially a very powerful tool. And as with all powerful tools, that means that if you do get this wrong, the outcome can potentially be bad.

It's also worth noting that DNS blocking doesn't actually remove content from the internet. All that it does is it makes it inaccessible to a specific set of users. This is a more limited alternative to actually suspending a domain. It is also, in some cases, more appropriate, or the only option that is available to a specific sort of community or set of users. As I've mentioned a few times now, it's also not 100% effective. There are a number of ways for people to circumvent DNS blocking, and we'll cover some of those in later slides. So next slide.

There we go. So how is DNS blocking actually accomplished? Here is a standard view of DNS resolution. And with DNS blocking on the recursive DNS server, the one in the middle, there is a filter list applied. When a user performs a DNS query, their query hits the recursive resolver. And if the domain is not on a block list, the DNS resolution just happens normally and the user is able to access the content. If the domain name is on a block list, instead of the user being able to access the information, something else happens with the DNS query. And, ideally, the user isn't able to access the query.

So when I say something else happens, there are a couple of different things that could happen. One of the more common, and potentially less harmful to the user, is the recursive resolver can respond with a response that says this domain name does not exist. So if somebody's trying to reach foo.bar.com, and foo.bar.com is on a block list, when the user tries to resolve that, the recursive resolver sends back a response saying, this domain name does not exist. The user then is unable to access foo.bar.com. Next slide.

Another thing that the recursive resolver can do is it can modify the answer from the authoritative server. And it can return a different IP address than what the authoritative server had intended. This is sometimes called redirection, and it is something which we see mandated relatively often. It is, I personally think, one of the least good options because it has a bunch of fairly bad security outcomes, which we will get to in a bit.

A third option, which initially to many people sounds like the best, is the recursive resolver can choose to simply ignore the query. If somebody tries to look up foo.bar.com, the recursive server could simply ignore this request and not answer at all. The downside to that is most people's computers, most users, have multiple recursive resolvers listed. And if the first one doesn't respond, the user's machine will simply ask the second or third or fourth. If that's recursive resolver isn't also performing the blocking, the user will still be able to resolve the name and get to the content which was supposed to be blocked. As we say down at the bottom of the slides, in italics, this only really affects the recursive resolvers where the blocking is happening. If users are using different resolvers, the blocking isn't going to be effective. Next.

So, domain blocking is in some ways similar to, but also distinctly different to domain suspension or domain seizure. DNS blocking makes the domain inaccessible to people using a specific set of recursive resolvers, whereas domain seizure removes the domain name from the internet completely, or changes the resolution. For anybody looking up the name. This domain suspension, or domain seizure, is more--

Excellent. Yay, technology. And now we've got an echo. Whose mic is on other than mine? Oh, well. So domain seizure ends up by going to the authoritative server or the registry or registrar and getting the-- Wow, the echo is annoying. There we go. Yay. So, domain seizure is usually used for things like botnet takedowns, or potentially something like trademark infringement or phishing

sites, where the information is updated in the authoritative DNS server or the registry or registrar. And this is a global operation. It affects everybody doing the resolution of the name.

This is a larger hammer in that it affects everybody using the domain name, but it's sometimes the only option which is available to somebody. For example, in a trademark infringement case, you can get potentially the domain seized or suspended, whereas you wouldn't be able to go along to every recursive resolver operator in the world and say, can you please add this name to your block list? How are we doing for time? I should probably speed this up some.

JONATHAN ZUCK: The Great Firewall of China is that doing what you're suggesting? Is it going all the way?

WARREN KUMARI: So, the question was, what does the Great Fall Firewall of China do? Basically, what it does is something similar to redirection. We're recorded twice. So the great firewall of China basically watches DNS queries that flow through it, if they're happening in the clear, and it inserts itself and replies with a different answer. So it's not actually doing what we're traditionally calling DNS blocking, in that it's not operating on the recursive. It's instead a man-in-the-middle attack, where it changes things on the fly.

So, motivations and examples. Why do people do this? There are a couple of different reasons. Probably the least controversial is

network operators perform DNS blocking on their recursive resolvers for security purposes. Oh, next slide. Sorry, I was changing slides on my laptop, but that obviously doesn't change it for everyone. Next one. Yay, that one. So, a significant number of recursive resolvers already do things like this. If the site is known to be hosting malware or some command and control thing for botnets, many recursive resolvers will perform this blocking. One of the well-known ones is, for example, CloudFlare runs 1.1.1.1. They also have a second IP address, 1.1.1.2, which users can choose to configure on their machine, and that will block access to known malware sites.

Next, another, slightly more controversial one, is, many enterprises have corporate policies that say things like employees should not browse social media sites like Facebook from their corporate machines or during work hours. And so they can implement this blocking by subscribing to a DNS block list and importing that into their recursive resolver or into their firewall. And then, if a user tries to resolve, for example, Facebook, the recursive resolver can block that resolution. And potentially redirect them somewhere else.

Another thing for this is certain organizations like schools, libraries, etc. Are required to block access to inappropriate content, adult content or gamemaster. Two examples of this. So, as I mentioned the earlier one, CloudFlare has 1.1.1.2, which block access to malware. They also provide 1.1.1.3, which blocks access to malware and adult content. So, if you're a parent and don't want

your kids accessing inappropriate content, you could potentially configure your local machines to use 1.1.1.3 as an example. Next.

And then this is one which gets a lot more of the press and is in some ways more controversial. This can be government-mandated blocking or court-mandated blocking, where a government or court instructs the recursive operators within their jurisdiction to block access to certain sets of domain names. This is known as one of the censorship mechanisms. Basically, content which is subversive or pedantically destabilizing, a government could mandate that all of their ISPs recursive operators within their country add a specific name to a block list or redirect the answers somewhere. There are a bunch of case studies in this paper. This is obviously just a short overview. Please actually read the paper. I know people probably won't, but please do. Next slide.

The SSAC is trying to remain technical in this paper, and we're not actually taking a stance on what counts as censorship or legality of specific cases. Again, please read the document for the disclaimers, etc. Next. And next. So, as I've said a number of times, DNS blocking is only partially effective. Down the bottom of the slide, we show a normal DNS resolution. The user sends their query to their ISP's recursive resolver. If the name is not on a block list, resolution works normally. If it is on a block list, the DNS resolution fails or is redirected somewhere else. Do we have a question?

JONATHAN ZUCK

Yeah, it's just a nit on the graphic. In the bottom example of the block domain, you still have an arrow going out there, but there's nothing that gets past the filter in that case, right? It just returns from the recursive server.

WARREN KUMARI

Yes, that's not necessarily a brilliant done graphic. It's supposed to show that if the name that's being looked up, if it's not on the block list, then that's the green option. If it is on the block list, that's the red option. It's not actually supposed to be showing where the user goes. I'm not sure if I actually answered your question, but yeah, the graphic is potentially confusing. That's supposed to be the outcome of a decision, not where the user goes from that. So this obviously all applies only if the user's query is actually hitting the ISP or the filtering resolver. If the user sends their query to a different resolver where there's no blocking, the blocking is not going to be effective. So, next slide.

How does this happen? The user can simply choose to use an alternative DNS server, not the one that is provided by their ISP, or not the one that's provided in-country. There are a bunch of ways of doing this. One of the better-known ones is public recursive resolvers. There are a number of these. Well-known ones are things like Google Public DNS, also known is 8888. CloudFlare runs one, Quad9 runs one. There are also government-sponsored recursive resolvers, such as the DNS4EU project.

In the, I think it was called Arab Springtime, there's a relatively well-known picture of, there was a bunch of DNS blocking happening in certain countries. I think the well-known picture is from Egypt, and people had sprayed 8888 on a bunch of buildings. So if you're Internet access was blocked and you wish to see content that was being blocked, you could just update your recursive resolvers to 8888. So this is something that's relatively easy to implement, especially for non-technical users. Jeff Houston's APNIC data shows that around 21% of worldwide users are currently using some of the well-known public recursive resolvers. Next slide.

If you've watched any streaming content in the last many years, you have likely seen a bunch of ads for various VPN services. NordVPN is a well-known one. Surfshark as another. And what these are designed to do is enhance privacy and anonymization for users. But what they also make very clear is they provide a means to bypass geographical restrictions. If you are traveling and want to watch a set of videos which are only available in your home country, you could install a VPN and VPN back to your home country. And now Geo-locked or geo-restricted content is available to you as if you were at home.

Another common case for this is if you are in your home country and want to watch media or content which is not available in your country, you can install a VPN and choose a different country to appear in. Now you look as though you're actually in that country. One of the things that this means is if there are DNS blocking or DNS

filtering happening within your location, you can simply install a VPN and bypass this.

And this is increasingly simple to do. There are ads all over most streaming services. If you look online, there are also VPNs being provided in many bits of customer premises equipment. Phones, similar. For example, Apple has Apple Private Relay. It's a simple thing to just toggle on. And suddenly, you look as though you're somewhere else. CloudFlare and a number of other organizations also provide basically a free VPN service to enhance your privacy, and they provide alternative DNS servers as well. Next slide.

Very similar and related to this are services like Tor. These are often marketed or provided to provide things like journalists and dissidents. A way to safely get out of their country. And this works by wrapping the user's data in multiple layers of protection, and each node along the path removes one. And so it's unclear to anybody who the user is, where they come from. It works really well, but obviously it provides another way to bypass DNS filtering because you can simply use resolvers not in your country. Next slide.

Another thing to keep in mind is there are things like DNSSEC validation. DNSSEC is specifically designed to detect and potentially mitigate people trying to fiddle with or modify DNS responses. This is exactly what DNS blocking does. It changes what the DNS response from the authoritative server was. We're

increasingly seeing that operating systems and users are having local validating resolvers.

A number of operating systems, like various Linux distributions, come with a built-in DNSSEC validating resolver. Increasingly, the home routers which you buy have DNSSEC validation in them. And so, if the blocking is applied on a specific set of the user's primary resolver, DNSSEC will detect that something bad has happened. The result will be bogus. And the user is likely to just apply the next resolver in the chain. If that one doesn't also have DNS filtering, they will be basically bypassing the filtering. Next slide.

I should probably speed up with my slides. The main thing to take away from this is that the success or effectiveness of DNS blocking is a matter of degree. It's in no way, a completely foolproof mechanism. All that it does is make sure that a specific set of users, using a specific set of recursive resolvers, will have their resolution blocked. But users are able to bypass that relatively easily. Install VPN, change your resolvers to others, etc. Next and next.

So the DNS has been one of the primary privacy leaking technologies for the last many years. We have increasingly trained users to use HTTPS sites, which encrypts the actual data that they're sending. We've increasingly improved the security and privacy of protocols, but the DNS has still been leaking where users are going. In order to help mitigate that, the IETF and other groups have been trying to protect the DNS. And one of the ways to do this is by encrypting the DNS lookups and DNS responses. There are a

number of flavors of this, DNS over TLS, DNS over HTTPS, DNS over QUIC. We're generally just calling them Dox at this point because there are many different transports.

But what this means is the DNS queries and responses are encrypted and often are indistinguishable from any other browsing data or any other internet data. This means that users can update their resolution system to use an encrypted resolver somewhere else, and this won't be detectable to filtering devices along the path. If the user doesn't want to use a blocking resolver, they can basically bypass the blocking, and this is invisible to the network. Next.

Another recent change is extended DNS error. This is RFC8914. And it was written to solve a number of different use cases, one of which happens to touch on DNS blocking. Traditionally, the DNS has provided almost no useful error information. If a domain resolution doesn't work, you would generally get back something like NxDomain, or, more commonly, server fail, which is basically something went wrong and I won't tell you what. This new RFC provides a way to annotate DNS messages and failures, and it has a bunch of different codes that can be applied.

One of these is Extended Error Code 16, which is basically censored. So, if a name is on a DNS block list, the recursive resolver can return, this domain doesn't exist. And the reason I'm saying that is I was compelled to by a court, or I was compelled to by an external party to sell this. They can also return different error messages or

annotations like this is on a block list which you have opted into. You have asked for this filtering to be applied, for example, 1.1.1.3. If somebody were to go to adult content, it could return you have opted into this filtering. Next and next.

As I said, DNS blocking is a tool. You need to understand its implications. If you perform blocking badly, there can be some bad unintended consequences. For example, if you're trying to block foo.bar.com and you accidentally block or unintentionally block the entire .com TLD, this ends badly for you. Let's check the next slide. Yep, next slide.

There is also potentially significant collateral damage if you choose the incorrect set of infrastructure to block. If you're trying to block foo.bar.com and you notice that it's hosted on, for example, Akamai, if you block the Akamai name servers, this will block all services which use Akamai, which is potentially a significant number. Again, this is just a high-level summary. Please see the case studies in the paper. Next.

When talking about the different types of blocking responses, I mentioned that the recursive resolver could return and X domain, which basically says the name doesn't exist. It could try and redirect the user somewhere else, or it could simply not respond. The redirecting the user somewhere else is something which we see recursive resolvers being compelled to do sometimes. And I think this is potentially one of the most dangerous options. It says if somebody tries to look up foo.bar.com, please redirect them to

this other IP address, usually where there is a web server providing some information as to why the user's access was blocked.

The danger to this is fairly much, by definition, where the domain is redirected to won't have a TLS certificate. So when the user is sent there, their browser will provide this pop up, saying attackers might be trying to steal your information. Do you want to proceed? And by doing this, we're training users to click through this important security mitigation. We've trained users to only go to HTTPS sites and not to click through these certs. This feels like by compelling recursive resolvers to use this answer, we're undoing all of those work. And we're training users to ignore TLS warnings. Next slide.

And in the interest of time, we will skip this and go to the recommendations. So next, there we go. The document only contains three recommendations. The one of the primary ones is if you're going to implement DNS blocking or mandate it, please make sure that you understand the implications of this, the strengths and weaknesses of the tools, and how to do it in a way that's not going to cause collateral damage. Next.

Recommendation 2 has a couple of sort of sub bullets. Basically, before doing DNS blocking, you should understand the implications of this. Will it actually be effective? You should have a clear policy about what to block and when, and potentially when to review these decisions and unblock if needed. Also, if you do implement DNS blocking, you should minimize the overblocking or

collateral damage, block only the most specific set of things needed to remove the content that you want to. And also, if you're implementing DNS blocking, only attempt to do this for users that are within your administrative or jurisdictional control.

And then next, which is almost my last slide. If you're a cursive DNS operator and you are returning any sort of blocked answer, you should use this RFC8914 extended DNS errors to annotate the answer in a way so that the user can understand why their access didn't work. If you don't do this, users are likely to become confused and will try changing things randomly, potentially break their machine, change their resolvers to ones which are not trustworthy or which potentially are sort of information collecting. And next.

Okay, we had a slide which had a pretty QR code linking to the document, but that's disappeared. I only have two or three minutes, so questions if any. Woohoo!

JONATHAN ZUCK I'll ask a question. Thinking in terms of how we began this conversation about the work begins with getting information out, who is your intended audience? Is it primarily those that operate recursive servers? Would it make sense for us to focus on things that have particular implications for end user confusion or habit changes or something, as you mentioned? I'm curious what the elevator pitch of this is.

WARREN KUMARI

Depends on who we mean by we here. For ALAC, I think one of the primary audiences are things like enterprises who mandate blocking towards their users. If you run an enterprise and you are suggesting DNS blocking as a way to make sure your employees don't access Facebook, these are the things you should keep in mind. Make sure that your employees are using your enterprise resolver. If you are doing blocking, probably you should not redirect the user somewhere else because you're just training your employees to ignore the TLS warnings.

From a more Essex audience or more global audiences, we were aiming this primarily at government regulators and policymakers who may decide for censorship or trademark infringement or blocking of objectionable content within their country. If they are going to be deploying this, what they need to keep in mind. Different audiences for different sets of people. I think from ALAC, make sure users understand if their DNS resolution doesn't work, potentially how they can figure out what caused that. If you are an ISP operator who's being compelled to do blocking, please annotate blocking responses with extended errors so users understand. And if you're an ISP doing this or an enterprise doing this voluntarily, these are the things to keep in mind. Only block what you need to, etc. I think Ram.

RAM MOHAN: If you could go to that slide that Warren was talking about where we're inadvertently training users to circumvent security, I think there is a relevant touchpoint.

WARREN KUMARI: Yep, that's probably slide 28. That one, whoa, we passed it. It's 28 on my number. Next one, next one. It's got to be that one. There we go. So yeah, if you are redirecting the user somewhere else, if they have tried to go to foo.bar.com and the DNS resolver has returned some different IP address and is sending them to an IP address that's not actually affiliated with bar.com, the web server that is running on that IP address won't be able to get a TLS certificate for foo.bar.com. TLS is designed specifically to prevent this sort of attack. And so the user will get a popup like this in their browser. They will be likely to become confused and try and figure out why can't I access foo.bar.com, and they will simply click through the warning message. Is that what you were getting at?

RAM MOHAN: Yeah, that's what I was referring to, Warren. Also, is it a normal experience that in many publicly accessible open Wi-Fi spots, airports and other public areas, there is some level of blocking or there's some things that happen that also prompt errors like this?

JONATHAN ZUCK: It certainly is on airplanes.

WARREN KUMARI

Yeah, so I think what Ram is referring to is captive portals. When you show up at a hotel or an airport lounge or a coffee shop, when you try and join their network, often they need to provide some sort of terms of service or acceptable use policy that you have to click through. The way that they generally do this is when you first connect, they hand out their own recursive resolvers. And any lookup that you try and do, they will intercept and send you to a web server under their control, which should pop up the like, hi, you're using Starbucks Wi-Fi, please click here to accept our acceptable use policy.

This used to be a larger problem, but is being mitigated somewhat these days, not completely effectively. Generally, when my phone, laptop, etc., joins a new Wi-Fi network, the network can send a captive portal response saying you're behind a captive portal. Here is the web server you should contact. Or the device, when it first connects, will open these sorts of connections in a sandbox browser environment. And so that does still happen. It is potentially happening less, but it's still, if it's not well implemented, can cause these sorts of pop-ups. Basically, what the captive portal is doing at this point is basically a DNS blocking or DNS redirection attack, but for a sort of more potentially more reasonable case. Did that actually answer the question or did I answer a different question?

RAM MOHAN It did, it did. And John?

JOHN LEVINE Hi, it's John Levine. This is a fairly long question, but I'll try to ask it in a short form, which is NIST in the U.S. And the standards organization has always had DNS recommendations and they have just revised them. And the new revision is completely different from the old ones. The old ones basically said, here's how to configure your name server so people don't hack it. The new one argues fairly strongly that enterprises should use DNS to firewall off stuff that is bad for their network and bad for their users. And honestly, they're not wrong.

But I think the reason why what they're saying is so different from what we're saying is that it depends so critically on your environment. If you're running the network in some company, then you're providing facilities to your employees to do stuff that's okay for the company. And if they make it difficult to go to Etsy and waste time buying stuff, that means you should do it on your phone at lunchtime.

So I think it is important to understand that our advice to governments is that if you're going to force blocking for everything in the country, you need to be super conservative about that. If you're in an organization, the block needs to match the goals of the organization. And typically, if you're at home, like 9999 actually does a fair amount of blocking to keep malware out. And they have a related one, which I think is 991010, which also blocks stuff that

is not family friendly. And they're quite upfront about this. If that's what you want, that's what you should get. So I think we need to balance like don't ever block because there's lots of good stuff that people should be able to see. On the other hand, saying no blocking at all just makes you look out of touch.

WARREN KUMARI I think that some of this is whether the user has opted into the blocking or not. If you work for an enterprise, for example, and they have some set of restrictions on what you're allowed to do, you have somewhat implicitly opted into I will follow the set of restrictions. It is also in things like an enterprise or a home network, the devices are likely managed. So in most corporate environments, if you bring your laptop or sorry, if your laptop is a corporate one, it's managed and the enterprise has forced you to use a specific set of resolvers, and you can't change them.

RAM MOHAN I see one, two, three questions.

WARREN KUMARI Let's go fast. I'll just move on.

UNKNOWN SPEAKER I'll go along. First, thank you for the report. Thank you in the sense that it draws the attention of the governments to DNS and not to really working blocking methods. So they do not harm the Internet

access for the end users in a way they could. That's fine. You point to the least intrusive way to implement blocking. That's fine. Unfortunately, at the moment, most recursive resolver software is not able to report that we are running in a censorship query. Hope that will change over the year. So thank you for not pointing out how to do real blocking. And I hope that you will never be. Thanks.

RAM MOHAN Robert.

ROBERT GUERRA Yeah, I just wanted to answer Jonathan's question and historical in terms of the first DNS blocking advisory that we did. I think the audience is, I would say, twofold. One for policymakers, and we're seeing increasingly in the space today, whether it's in the U.S. or other places as well, too. Increasingly, the type of content that is wanting to be blocked or removed from the Internet is occurring in a way that is not transparent. And so if it's going to be done through DNS blocking of some kind, at least my view, and I think the SSAC as well, is to be as open and transparent as possible. So if they're going to do that, what are the consequences? How are they going to impact the system?

And so part of the DNS blocking report that we've done or the update is that given developments of how the Internet's being used, here's some of the consequences if you do block what's going to happen. And I think for users, what Warren mentioned about the

captive portal and other things, too, is that for users who then want to get around it or want to know how blocking takes place, it's an educational resource for them as well, too. So I think that's particularly important.

And I think what needs to recognize is that, you know, particularly in some parts of the world, a lot of content is disappearing. And so it's really important for users to know why that's occurring in place if they want to access it in some way or another, too. So it's educational, and it's written in the language also, so it's not just for the technical experts, but as the first blocking report was done as well, to inform some of the collateral consequences. And if you're going to block, what are some of the issues? And then it's better that there be a fee for information. Thank you.

RAM MOHAN Thanks, Robert. Matthias, last comment on this topic?

MATTHIAS HUDOBNIK Thanks a lot. Matthias Hudobnik speaking. I would just ask Warren if he can maybe say in a nutshell two, three practical points from the end-user perspective, what they can do and what they can take out of it. I mean, I know it depends on the blocking technique, but in general, mind your audience. We do not have a technical audience here. What is really, let's say, two, three points what you would recommend or, let's say, what do you think is meaningful? Thank you.

WARREN KUMARI

So I think the first thing is to decide whether or not you want a specific set of blocking. If you have decided that you want adult content or malware to be blocked, you could potentially update your recursive resolver to a service which provides that. There are a number of them. But if you are a user that is subject to blocking, which you would like to avoid, things that you could do is update your recursive resolver to be one of the public resolvers. Or depending on where you are, you could also just do something like sign up for a VPN service and VPN your traffic through something. It depends, again, on local laws and how much you want to do.

I will note that with a VPN, your traffic appears somewhere else. And all of your traffic is passing through the VPN. So you need to choose a VPN provider that's trustworthy. But also, your web browsing is likely to be slower because your traffic is going potentially a much longer distance. But generally, change your resolver VPN.

JONATHAN ZUCK

Is there somebody?

RAM MOHAN

Great. Thank you. So let's go to the next update for you. Maarten is here. He is one of the cool work party chairs for the Free and Open Source Software Work Party. Over to you.

MAARTEN AERTSEN

Yeah. So within SSAC, we're working on a report that is not finished yet. But it's on the use of free and open source software in the global DNS and domain registration infrastructure. So the motivation for this is that some of us noticed that in policy development processes around regulating software or regulating infrastructure, the use of free and open source software was basically an afterthought. And we knew that this is being used widely. So while finding for papers to cite, we came up basically empty.

So this paper tries to document the use of free and open source software in the DNS and document its characteristics. And then the idea is to make the reliance visible and to be a resource for policymakers to make effective policy without considering this as an afterthought. So for those that are maybe not as familiar with free and open source software, this is a software which allows anyone to use the software, but also to study, modify, and distribute it, which ended up being quite aligned with early culture on the internet for software such as a DNS software. Next slide, please.

So we're investigating four areas in this upcoming report. Two of them are basically like a survey activity, and two of them are more of analysis and policy considerations. So the first two map the usage of free and open source software across two sets of infrastructure. So the first one is the infrastructure that is focused

on registering domain names, which this audience is very familiar with. We specifically look at registries. And then the second part is the DNS, where we look at the usage of FOSS in name servers and resolver operators. So that's our infrastructure landscape survey.

The second part to this is the survey of the characteristics of free and open source software. So how is this different from physical goods and proprietary software? Because it turns out it is very different. And we then contrast this type of software generally to the open source software, which is very popular in the DNS space. The third and fourth are then to list and clarify a number of common and actually quite logical assumptions that people make when thinking about this, and to actually clarify that these are misplaced. So the fourth is document the risk factors we deem relevant when making policy or when regulating. So this applies both to development of this software and to operations of this software, but also to regulatory approaches. Next slide, please.

So I'm not giving you answers today because I'm basically teasing what we're working on and we hope to give you a full presentation in Muscat, which also will include guidelines. So looking forward, we feel that continued use is a strength, but there are risks. And in our guidelines, we will offer our view on this matter. I think this is where I will stop today, unless we have more time, then I can dive into each of the pies and I'm looking at the chairs to hear what you want.

RAM MOHAN Why don't we pause for questions and perhaps we finish all the presentations. There's more time we can come back to you, Maarten.

MAARTEN AERTSEN Yeah, I think that makes sense.

EUNICE ALEJANDRA PÉREZ COELLO Thank you. My name is Eunice. Does this software work with any operating system?

MAARTEN AERTSEN Thank you and apologies for the difficulty with me using the translation device. So the question was asked whether this software is available for any operating system. So most of this software is available for multiple operating systems. Operators tend to make use of a certain set of operating systems and this software is available for all of them. So while I can say, yes, it's available for all, for those that are actually used in practice, the software is tailored to them. Yeah, yeah. The report usually mostly looks at the top layer. So the application that provide the DNS service, we don't look below, but we do know that below there's a lot of this stuff as well. Yeah. Thank you.

RAM MOHAN Other questions? No. Okay. Not yet. We hear the punchline. Yeah. Not yet. Okay, let's continue with the rest of our presentations and

then if there are other questions, we can come back, Maarten, or any other topics that are here. Let's go to the next slide. The next presentation really. There you go.

So the SAC is also looking into, we've started up a work party that is looking at the responsible integration into the DNS ecosystem. In other words, blockchain identifiers among others and their impact into the DNS ecosystem. You see what's on the slide there. We're trying to do some base groundwork on what does integration mean and then what are the risks when you try those kinds of integrations? Because it's relatively speaking, fairly open-ended right now and different operators define integration in their own way. And they're offering, in some cases, they're making either promises or they're asserting certain types of behavior that may or may not actually be guaranteed or something that is deliverable in some kind of a consistent way.

When you look at a namespace that is not the DNS and you want to integrate that non-DNS namespace into the DNS namespace but have some kind of interlink and have some kind of constant management between one and the other, there's bound to be problems and issues. And really the underlying piece of this, and we've spoken about this at several plenaries and other forums inside of ICANN before.

The fundamental thing is, as with all technology, spaces evolve, namespaces evolve, new namespaces come through, there are new places for identity to be expressed. That's not the issue. The

issue is that if you look at the domain namespace, it has a certain set of nomenclatures and norms and rules and expected behavior and there's a level of trust that has been hard fought and established over a long period of time. Any new space that comes through is clearly going to try to leverage that trust and repurpose the names. So you call it a resolver even though it's not a DNS resolver, you call it a domain name even though it is not a domain name in the DNS space. So you repurpose the words because the words have established meaning, trust and knowledge for end users, for enterprises all over the world.

And one of the risks is that when you repurpose not just the names, but then you take some parts of the functionality that you know works in the DNS space and then you marry it to other pieces of functionality that is unique to other namespaces and you merge them together, now you start laying the groundwork for pretty good user confusion. What is a domain name? What is my identity? Will it really work? Why is this not working on my browser? Can I send an email to it? All the things, all the behaviors that people expect to just work as is in one namespace, we want to make sure that there is some clarity in that area among other things. But this is still in the early stages.

The next slide, we're actually right now, while we have the long-term focus, in the short term right now, we're pretty focused on impact to the next round of gTLDs. And we noticed that ICANN has put out the applicant guidebook. And this is several months ago, sometime last year, the SSAC published the name collision analysis

project with a set of recommendations. And we're pleased that the Board took it seriously, adopted those recommendations, and ICANN Org is implementing those recommendations as a part of the next round. But we're intending to apply our analytical framework to the next round because we think that there is a possibility of user harm that comes from collisions with alternate naming systems that we had not explicitly addressed in our prior report.

So that's kind of a key focus and this work party will be addressing that. The work party actually has a session at this meeting where we're going to be delving into those topics. So those of you who are interested, please feel free to come in and listen to the deliberations of the SSAC. Questions?

JONATHAN ZUCK I'll ask one quickly. First of all, name collisions is near and dear to the At-Large community. I think it's one of the first kind of collaborations between us, because I think we helped to raise alarm about name collisions when your first report came out to kind of relative PhD level anonymity. So I think this is something near and dear to us. Do you think you'll come out in time with a recommended methodology for doing this? Or will it just be another kind of cautionary tale that this could happen or that could happen and everyone should be careful?

Do you think there's going to be a right way? Because there's already been some promises made like .APE where they're trying to

do a very smooth integration by getting .APE as a community priority evaluation. I don't want you to bad mouth anybody, but I mean are you going to be able to be far enough down this road that there is like a right and a wrong way to do this?

RAM MOHAN Warren?

WARREN KUMARI So from sort of a personal thing, one of my things is mainly making sure that people are aware and there's a plan in place ahead of time. During the last round, there was the drama around Home Corp and Mail and a fair bit of it was there wasn't a process and plan and procedures in place ahead of time. And so I think one of the things is making sure that we are prepared for when this happens.

Everybody understands what the process is. Everybody understands what the policies are. What the actual policies are is not necessarily secondary importance, but basically be ready for this. Make sure that everybody is informed and that if this actually happens, if there are multiple applications coming from different blockchain providers, is there any sort of priority? Do you take into account the fact that there's an existing set of users or do you explicitly not basically make sure that the process is understood and prepared for ahead of time?

RAM MOHAN Yeah, so Jonathan, I don't want to preempt the work of the work party, but I think it's unlikely that we're going to come and say these things are good, these things are bad. I think we're much more likely to say there'll be dragons.

JONATHAN ZUCK Yeah, because the new round isn't that far away anymore.

RAM MOHAN Justine?

JUSTINE CHEW Thanks, Justine, for the record. So I think this is a very interesting thing and I think it's absolutely necessary. I can tell you from my experience in attending the APAC DNS forum, which was back-to-back with the CP Summit, and you were there, right? We have clear expressions, not even of interest, of confidence that these Web3 players are going to come in and get names that they want, regardless of what the rules are. That is the expression that I took as a takeaway from being there.

So this is of very much concern. And they are probably going to be using the grounds that they already have use of that string in their own namespace as a solid foundation for getting the string. Regardless of whether anyone else applies for that string. So I think this is very, very important. We would be very appreciative if there

were some things, even coming out now or before the application round begins, that can alert the board to these ramifications.

Sorry, just to say that the tactics that these Web3 players are using is a lot of jargon which people don't understand. So they're coming to say things like, you know, Web2 should learn from Web3. I'm like, what the hell is Web2, you know? There's only the web and you Web3. And a lot of things that the Web3 prescribe is not really practical for the entire internet. So nobody really understands these things. They're just going around throwing names like they're God gives to the earth or something. So it's very concerning.

RAM MOHAN Thanks. And we're aware that the clock is ticking and the next round is right around the corner, which is why the work party is reorienting its focus to specifically look at the applicant guidebook and providing specific advice or at least our comments that would then be considered by ICANN Org as it looks to revise the applicant guidebook taking into consideration public comments. Because that's kind of the avenue open for us. We've concluded our work with Name Collision and we've realized that this part was a piece that was the hanging chad. All right, next slide.

Let's go to the last of our work parties that we're working on. This is DNSSEC operational considerations. We've just begun the work in this area. The fundamental issue that we are trying to solve or trying to address with this work party is DNSSEC itself, as you know,

has been around for a while now, but it's yet to become an overnight success. And we're partly spending some of our time looking at why it's not more widely used.

Are there ways to resolve that? And identifying what are the hurdles, what are the challenges that exist? And are there practical solutions that perhaps either should be built or already exist and haven't been understood well enough. So that's, again, the high level of what we're trying to do. This work party is also meeting here at this meeting in public, and that's on the ICANN schedule. You should be able to come and listen in to our deliberations on this work party.

That brings us to the next slide, which, Jonathan, you and I were talking about this at the last meeting and meetings prior to that. And if you please go to the prior slide still. Safer cyber was something that we came up with together, and I think it's been off to a good start. You know, there's already something that you've built and is now part of in a curriculum and a dissemination everywhere. And that's very encouraging for us. So we are coming to you, and we want to chat with you and open up a dialogue now and in the future on what safer cyber is not just one thing.

So if you go to the next slide, we think we're thinking there might be a variety of topics that we could work on and really want to open it up to you here in the room and online on which of these things may be interesting, may be relevant, and timely for your audiences. For the first two topics, we have a couple of documents that we

have written that are not directly aimed at the end-user audience, but there is interesting material there. You know, SAC40 and SAC115 are both documents. We can derive material from there for the first two proposed topics. But there are others where we've written strong reports that could be leveraged. So I leave it as an open topic here on, what strikes your interest and how we may perhaps progress this further?

JONATHAN ZUCK

I'm open the floor to everyone. First reaction is positive, and this new ACES group that Amrita is chairing, I think, will be the conduit for these conversations. And then we have a couple of folks spies within the SSAC, that hopefully come back with some action items for us and for the group on these various issues.

But I think we'd be very interested in a pipeline of trying to get some of this information out. And as we build alliances, stronger relationships with our ALSs, with libraries, as we talked about, and things, I think there's no reason this wouldn't be an ongoing pipeline of information for folks. And maybe customizing some of the audiences for some of these things and looking at different means of getting information out, like you said, in animations or things like that, that sometimes require going out and finding funding and things that make it more difficult.

But as far as what we can produce ourselves, we're on board to get some kind of pipeline going. But let's figure out, let's look at the grant program as well, as a mechanism for taking the work, the

amazing work that the SSAC is doing and repackaging it in a way that's more palatable. That money is out there, and I think it's a great use for it. So let's have an expansive conversation about it that we won't have in two minutes, but I think certainly, I think it's a great idea.

Any questions or comments on that for Ram from our crowd here? I don't have the participants up to see if their hands are up online. Is there anybody here? I think we're out of time. We just enjoy ourselves so much here, the ALAC and SSAC, that we run out of time. We need to appeal to ICANN to allow us to have our own retreat or something, where we just have three days devoted to how to make end users safer online, and what are the tools and technologies and partnerships that we could form to do that. So let's get on that in the next budget cycle. We'll see what they're up for.

But beyond that, thanks again for coming to visit with us. We always enjoy it. We always learn something. And as you say, let's keep the communication going intersessionally as well. Thank you.

RAM MOHAN Thank you so much.

[END OF TRANSCRIPTION]