
SINGAPORE – SSAC Public Meeting
Thursday, March 27th 2014 – 08:00 to 09:00
ICANN – Singapore, Singapore

PATRICK FALSTROM:

...of the IANA function that will be presented this afternoon at the open session. If we look at our publications per category, we of course produced most reports lately on DNS security, but we also produce reports related to DNS abuse. And now with SAC 65 and also the [inaudible] work part of that, we are probably going to increase the amount of work we do in that area.

Internationalize domain name and also WHOIS. WHOIS, which of course is a longstanding item and Jim Galvin here to my right is a person that inside SSAC is trying to [check off] everything which has to do with WHOIS registration data and such items. So if you have any questions in that area, please talk to Jim. Is there any questions on this general description of SSAC?

So let's move on to then SAC 64, Advisory on DNS Search List Processing. In this report, SSAC examines how current operating systems and applications process search lists. What we are looking at is whether various applications do the same thing, because if there are discrepancies between how search lists are applied, then of course when you type in or when you use the same thing that looks like a domain name in two different applications, two different operating systems will result in different DNS lookups. And because of that, might end up resulting in different results.

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

So we look at security and stability implications with both different search lists, behaviors, and also the fact that the search list behaviors are different.

We also propose a straw man to improve the processing. So if we look at what search list processing is, it is that if you have a search list of somedomain1.com, somedomain2.com and a user type system into the browser search box, the operating system would try system.somedomain1.com, system.somedomain2.com and system in some order.

Some even try www.system.somedomain1.com and www.system.somedomain2.com and www.system. It's actually quite bad.

So we looked at the fact that there's non-standardization of search list, which we have categorized the various behaviors, and they range everything from not applying search list behavior at all to [inaudible] www to the regional [inaudible] or the string or some kind of string that is used as a search term to the default search engine.

Once again, the fact that this is non-standardized is not very good, specifically in combination of introduction of new TLDs, which I will come back to.

Given that this is happening, we have been looking at the issue with query leakage and privacy issues, and not only is it the case that the query itself might go to places and resolvers where the query would not go, it's also the case that if it is the case that the query actually resulted in a response, in that case, applications might try to connect to



whatever the [or] data was in that response to that IP address or similar, depending on what resource [record type] was queried for.

This, in turn, might imply that the application actually is leaking data to whoever intentionally or unintentionally did set up something on that specific domain name.

It's important to remember that the protocols of use on the Internet are not only HTTP, and depending on what protocol it is, the protocols do send different kinds of data before it even discovers that it is connected to the wrong host.

So even in the first sort of batch of UDP or TCP data it might be the case that the client is sending which one could say it's too much data. For example, in HTTP, even the first request might include cookies, which in turn include privacy information.

So how come this is so difficult at this specific moment? Well, of course it has to do with the fact that we now introduce new top level domains. And we have already before been looking at the DITL data that queries to the root servers, which is one of many different kinds of data sets one could look at. And we are updated the data that we have in SSAC that we fetch from the DITL and we see that, in 2013, DITL data, which is in red, there's a large increase in queries for top level domains which are not yet allocated, for example, .home and .corp, which are two of the applications.

If we look at the logarithmic Y-axis instead, which is better for some kind of – depending on what data it would like to have out of the graph, the previous one was linear. Y-axis, this is logarithmic. So the actual



problem has to do with the fact that search list processing starts with whatever domain name was typed in or configured or what people were looking at for – www.corp, this is one way of doing it.

And the whole search list processing relies on getting back an [X] domain, which means that the data queried for does not exist in DNS.

So this is one example where it starts with www.corp and then it adds whatever it is in the search list and tries various different kind of queries until it queries for www.corp.example.com even though what was typed in in this case was www.corp.

If you have a look at the first line, and then you see that is the query and the second line is the response, you see that if it is the case that we actually allocate the TLD .corp and someone registers the domain name “www” and in that domain have an A-record for www.corp. Then the response on line number two will no longer be an NXDOMAIN. It will instead return an IP address. That IP address will stop the search list processing and the IP address is handed over to the application, which is going to do whatever it’s supposed to do. That’s the simplest case of explanation, what kind of problem might occur.

So the straw man proposal is, first of all, that the search list should not be generated automatically. There are some cases in operating systems and applications where the host name of whatever is connected to the Internet, the box itself, is the host name including the domain name of the host name, the fully-qualified domain name of the local host. From that, a search list is derived and we recommend against that.



If search lists are to be used at all, that should be explicitly configured and this includes DHCP servers and whatnot.

And then we have two rule, which are complementary. We say that unqualified single level domain names are never to be queried directly. And this is something that is continuation of earlier SSAC advice that there should not be any A records or such at the apex of a TLD.

So if it is the case you have a token that is a single label domain name, one should always apply search lists and never issue a DNS query with just that single label.

On the other hand, the last rule, if it is the case that you have a multi-label domain name, then search lists should never be applied.

So I would propose straw man proposal is that given a domain name it should be known and well-specified whether search lists should be applied or not to that domain name.

So the recommendations include that SSAC invites SO and ACs, IETF and DNS operations community to consider our straw man proposal. We recommend ICANN staff to work with DNS community and the IETF to encourage standardization of search list processing behavior.

And what we mean is the existing RFC that talks about search list processing is (a) not clear enough and (b)—

[audio cuts out]



MERIKE KAE0:

...critical to the stability and security of the entire Internet. And as I mentioned before, the fact that you have increasing bandwidth, you have hosting providers with hundreds of thousands of hosts that an attacker, if it gets to compromise the hosting provider. I mean, this is a really bad problem and it's just growing.

So we explored several critical design and deployment issues and made quite a few recommendations to both ICANN and also the operators of the Internet infrastructure and manufacturers to take specific action.

And so as I go through the recommendations, I do want to clarify that Recommendation Number 1 was specifically to the ICANN Board as to what ICANN could potentially do to help with this problem. The other five recommendations were more targeted to the ICANN multi-stakeholder constituency, because it deals more with what we recommend manufacturers do with some of these devices that are problematic, and also network operators.

And this is not something that ICANN specifically has influence over, but we did feel that we should make these recommendations so that the multi-stakeholder community can also help alleviate some of these issues.

So the SSAC specifically recommends in Number 1 that ICANN should help facilitate an Internet-wide community effort to reduce the number of open resolvers and also networks that allow network spoofing. And this effort should involve measurement efforts and outreach, because one of the problems has been trying to determine how many networks globally actually do allow spoofing so that you can raise awareness that



you're causing the problem; could you please help solve this by creating – using mitigation techniques that are well-known?

Recommendation Number 2 is that all network operators should take immediate steps to prevent network address spoofing. And so this speaks to putting in mitigation techniques that deal with source address validation. One of my colleagues had come up with a great acronym as we were trying to figure out, "Well, what do we do?" It's like SAVE the Internet – Source Address Validation Everywhere – which is actually not bad.

Recommendation Number 3, that recursive DNS server operators should take immediate steps to secure open recursive DNS servers. I do want to also mention that this is unmanaged open recursive DNS resolvers, because there are some companies that provide these services, but they're also cognizant and have mitigation techniques for in case you have any of these amplification attacks.

But there's many devices that don't realize they run an unmanaged open recursive resolver, and basically they allow queries from anybody from anywhere, which is part of the problem. You should restrict where the queries come from in some environments.

Recommendation Number 4, that authoritative DNS server operators should support efforts to investigate authoritative response rate limiting (RRL). Some people have found that rate limiting actually helps with the problem.

Recommendation Number 5 is that DNS server operators should put in place operational processes to ensure that their DNS software is regular



updated and communicate with their software vendors to keep up to date with the latest developments.

For those of you that deal with any kind of other security issues, this is just updating your software and patch management. And we found that a lot of people are actually still running very, very old versions of some BIND software and we're kind of surprised that people don't update their DNS software as much as you would think.

The last one is personally what I feel is quite critical. This is that manufacturers and/or configurators of customer premise equipment – so CPEs, your home routers – should take immediate steps to secure these devices and ensure that they're field upgradable.

One of the problems with these devices is that people buy them, they pay 50 euros or 50 dollars, whatever it is, and then you plug it in and you just let it run for three to four years until you have to buy a new one for whatever reason.

The problem is that many of these devices have security issues. Many of these, as a default, run as open recursive DNS servers. So one of the reasons that we put this recommendation in here is that in terms of the multi-stakeholder community that ICANN has, we really want to raise the awareness that the manufacturers have to fix their equipment, and anything new that comes out, even with the Internet of things, that will have some kind of DNS capability to make sure that they actually do not run as open recursive DNS servers and are software upgradable. That's actually a key point.

So that's it and I'll open this up for questions.



PATRIK FALSTROM: So we'll start with questions related to this report, and then we'll open up to any kind of questions related to SSAC work, or maybe others as well. Well, you can try. A question from the chat room I understand.

UNIDENTIFIED FEMALE: Yes. This is from Steve Metalitz. He asks: "In SAC 65 section 4.2 it is noted that some recursive DNS servers to protect against abuse use the tactic of 'not answering for common amplification queries' and 'completely dropping requests for domains that have been identified as part of DNS amplification attacks.' Why do these tactics not bring the same risks e.g. to DNSSEC that SSAC has observed with regard to DNS blocking?"

PATRIK FALSTROM: It is actually very difficult to hear because the speakers do not hit us. What section was it?

UNIDENTIFIED FEMALE: Section 4.2. Do you want me to repeat the question? I can show it to you. That might be easier.

MERIKE KAE0: If I understand the question correctly, it deals with why do you do rate limiting rather than just blocking and dropping all the queries from a given domain? And the problem is that you could also be blocking legitimate queries. So even when you're rate limiting, the problem is



that you could be dropping legitimate queries, but it's thought that that was much better than actually dropping any legitimate queries. Did you want to comment?

UNIDENTIFIED MALE: If I may ask a question. This is not a technical question.

PATRIK FALSTROM: You have to speak much, much closer to the microphone because we don't hear you.

UNIDENTIFIED MALE: Okay. This is not a technical question. Urging good behavior is one approach to solving these problem. But what economic regulatory or other incentives are there that will encourage our force, to use the two extremes, appropriate behavior?

MERIKE KAE0: Quite frankly, it's a difficult one, because for the people that need to be putting some of these mitigation techniques in place, it doesn't help them. It helps others. However, if people don't do it collectively, you may find your business or your environment at some point being susceptible, which means that you cannot access the Internet because you have just too much traffic. So it's a matter of people understanding that, at one point in time, this may hit the also and it's really collectively best for the entire stability and security of the Internet to put in these practices.



From a regulatory environment, I think it's very, very difficult. There's one country that I had talked to – Finland. It's kind of interesting because I was talking to the Finnish network operators and I said, "How come you're so good about doing these anti-spoof filters?" They said, "Well, because the regulatory body highly recommends it and we just do it." But what is the enforcement?

EDWARD LEWIS: I actually want to follow up to his question.

PATRIK FALSTROM: Please, please. We have interpreters. You must speak much, much slower.

EDWARD LEWIS: Every time you recommend more controls to the operations of the protocol, you're going to raise the cost of the operators. So doing this is going to mean more cost to those who are running it today. That's a disincentive to add these controls. You have to add a positive reason to go ahead and this. To underscore what the previous speaker just said, I think it's very important.

MERIKE KAE0: I agree. And this is why we want to raise the awareness, so that we collectively can figure out how do you actually get to the positive incentives, because clearly over the last decade, we haven't been very good at that.



UNIDENTIFIED MALE: So I'm going to go back to the technical question. At one point, rate limiting was mentioned in the recommendations. Microsoft being those DNS servers. [inaudible] from Microsoft, by the way. I should have probably introduced myself first.

It does not currently support rate limiting, but will be supporting rate limiting very soon. The real question I have is what other recommendations you can make to the software providers, the people that actually deliver the software at this level?

MERIKE KAE0: At this point, as SSAC and from the report, I can't speak for SSAC but I can talk to you offline about that.

UNIDENTIFIED FEMALE: We have two other questions from the chat room. One is a follow-up point from Steve Metalitz. He says: "My question was not about rate limiting Recommendation 4, but about limiting or blocking known abuse queries." And that's Recommendation 3C.

PATRIK FALSTROM: In reality, I think people, for example, like SSAC [inaudible] is doing combination of rate limiting and blocking queries. As Merike said before, the [manual attacks] we see are actually on legitimate sources and legitimate queries, which means that rate limiting is the proper way of doing – of trying to understand what subset of the queries coming in should you rate limit and which ones should you not.



That said, there are illegitimate queries as well that one can drop. So I don't think it's black or white in one or the other, but a combination of the two.

UNIDENTIFIED FEMALE: And the other question is from [Ser]. Sorry, I seem to have lost it here. Hang on. It seems to jump around. [Ser] asks: "On Recommendation 5, why did we not include on the recommendation that DNS software must be automatically updated?"

MERIKE KAEO: I think from an automatic update point of view, I don't think it's necessary that's something that we would want to recommend because while we would highly encourage people to update software, I think most of us that [are in] uncertain environments know that sometimes a company has processes in place in terms of when to upgrade and do some kind of testing.

And so we didn't want to recommend that. We wanted to have the capability to update the software, to actually recommend and have people require that it be automatic was not something we wanted to actually require.

PAUL EBERSMAN: Paul Ebersman, Comcast. A couple of comments. One of them might be a gentleman who was asking about [RL and RPZ] had mentioned DNSSEC as well. So I will point out that RPZ is one effective tool, depending on



your environment. But one of the issues with it is, because you are you lying in DNS, you do break DNSSEC. It will not validate.

The other comment I'd make is that we've actually been looking at quite a bit with the CPE problem and there's also some very good research being done by [Jim Geddis] and [Dave Tott] and one of the issues is not just that people keep software for years, but there are vendors currently shipping today in stuff you can buy at Best Buy versions of DNS [mask] and NTP that are known broken and are years old.

MERIKE KAE0: Yeah. And I'll make a comment to that also. For those of you that aren't aware, there's a project called the OpenRecursiveResolverProject.org and it's identified over 30 million open recursive resolvers. So which ones of these are in a managed environment and which ones aren't, that research still needs to be done. But 30 million, okay, that's a lot.

PATRIK FALSTROM: More questions?

UNIDENTIFIED FEMALE: This question is in the chat room from Ralph. He asks: "What is being done to ensure that old versions of DNS software get updated? Isn't this a major hole? Who do you think should work on resolving this issue?"

PATRIK FALSTROM: Well I think all of us know that ensuring people upgrade software is something that doesn't really work. As was pointed out earlier, the



carrot is probably more effective than the whip here. So I think people will upgrade software when they want to have a new version of the browser. They can watch new video or whatever and that will result in upgrade of the software and the operating system, and they get the new features.

But we see the discussions going on about old version of operating systems, and without pointing any fingers, specifically for example, there's a big discussion going on on Microsoft XP at the moment where many people think that Windows XP is really, really old, but it's still out there. So it's really, really hard. Microsoft and others are working really, really, really hard to get people to do that upgrade, but it's just difficult.

But as we point out in this report, we don't only talk about devices that people actually have on their laptops and others where there might be open resolvers.

As Merike pointed out and some other people here in the audience have seen that, for example, all of these USB sticks or cell phones and other kinds of things. I've been working personally with a case where a USB stick for a laptop included the resolver that responded to queries also on the upstream interface, and that device is not upgradable and a telco, a mobile provider, in Sweden distributed 150,000 of those before they detected this problem.

MERIKE KAE0:

I'm also concerned about all of the home devices that actually are now Internet capable. I mean, I have at home waiting for me light bulbs that I can control using my iPhone. You can change colors just using your



iPhone. I also bought a dryer six months ago because my old one died. I bought the one that I could check via the Internet, just because it was interesting. But it's concerning me at this point because I'm going, "Well, if it gets access to the Internet, I'm sure it runs DNS and I don't have time to check out its DNS implementation." But it would be interesting to see, are these all part of the problem?

ALEX DEACON: Hi, my name is Alex Deacon. I just want to follow up on the question around rate limiting, because I don't think I quite understood it. Can you confirm that rate limiting will in fact [break] DNSSEC? Is that what I heard?

PATRIK FALSTROM: No. RPZ, okay.

ALEX DEACON: So maybe my question is: what is the effect of rate limiting on DNSSEC, if any? None?

PATRICK FALSTROM: Yep.

UNIDENTIFIED MALE: Rate limiting doesn't do anything. It can [inaudible] for DNSSEC. It's just basically the rate how the [inaudible] are going. It doesn't interfere with the [count], etc.



UNIDENTIFIED FEMALE: This is another question from Ralph in the chat room. He says: “I’m talking about BIND in particular. With a terribly broken and compromised installations in many parts of the world, it causes real problems in many countries. There seems to be no [inaudible] on that aspect. OS, etc., have vendors working on it, but DNS software seems to be an orphan. Isn’t there something ICANN should do?”

MERIKE KAEO: Since he asked “Is there something ICANN should do?” I’m in a little bit of a conundrum, because I don’t want to answer for ICANN, but I want to answer as an SSAC member for security and stability. I think that ICANN has been doing a great job in terms of outreach to the overall community, and I think part of that outreach maybe could include just awareness that if you’re running older versions of software and actually talking about what an old version of software is, because some people just don’t know. Somebody installed the DNS maybe five years ago. The Internet is working. What’s the problem?

So I think it’s just an awareness that best current practice you should be running version XYZ or higher, and maybe an awareness program would be useful.

DENNIS JENNINGS: Dennis Jennings again. Another non-technical question. Has SSAC the resources it needs given its increasing workload, or it seems to me its increasing workload? And, in particular, if SSAC needs to do some



foresight activities to anticipate problems, does SSAC need more resources?

PATRIK FALSTROM:

Answering as the chair of SSAC, our resources are of course something that we are discussing both internally and also with the ICANN staff – staff support and also of course the Finance Committee that there is actually a meeting just after this.

First of all, we do believe that we have good support from ICANN. Let me start there. We are careful with our requests, and most of the requests I must say that the Board approves, and also [approved] by ICANN, which means that to some degree, everything is okay. Of course there's always some things that you can discuss with them, but in general terms, the situation for SSAC is good.

Now, could we increase our output? That's a slightly different question. Let me go into that. SSAC consists of volunteers. All of us donate our time. All of us have a day job. We do have at the moment three individuals from ICANN that do staff support for SSAC, and each one of them are able to support approximately two work parties each, which means that we can have six work parties ongoing in parallel.

It takes some time to create some output and create these reports, and there are a couple of questions here. First of all, could we, with more staff support, be able to produce more results or have more of these groups in parallel?

We do have a discussion with ICANN that with a little bit more staff support, we might be able to move these six parallel groups forward a



little bit faster. And specifically, our staff support will not be as tired when moving into [last call] on these reports, so to offload them, yes it would be good to have more staff support – at least in the number of hours, maybe not in the number of individuals.

But we do not believe that we can have more work parties run in parallel. That would be difficult, because we have these 40 volunteers. So I think it's difficult for us to produce more than what we currently produce, which is about eight reports a year. And remember that if we remove the holiday seasons, that's about one report every month which I think is pretty okay for a volunteer group that tries to produce high-quality reports. That's really important to not forget that.

Then there is one last [factor]. That has to do with the fact that we just have volunteers. We don't have any real staff that, for example, can do investigations or research. That's why you see we do recommendations where we recommend ICANN to issue studies and similar things.

Then ICANN is the one that launches studies, and I do believe that you personally have been involved in some of them, like [ask] the community [to do] the studies. We from SSAC do believe that that is actually not a very bad idea because it makes it possible when doing the study to sort of gather all different kinds of people, individual organizations that are suited for that task to do more research, outreach and do the study itself. Then SSAC has an ability to comment on the result of that study.

So I think we are pretty happy where we are at the moment. Just because of the way SSAC is constituted, we might be able to produce a little bit more. But on the other hand, I don't really know whether that



would be so helpful either, because if we look at the number of reports and recommendations that we issue, there's pretty slow uptake on our recommendations as well. So if we increase the number of recommendations, I don't think the world would get much better.

Please?

MARILYN CADE:

Thank you. My name is Marilyn Cade. I actually have an idea I'd like to explore with you about how to make the world a little better through the use of your reports.

I happen to have the benefit of having extensive technical training when I worked for a high-tech company a number of years ago. It's not current training, but I allows me to take the veneer approach to understanding these issues. That is, understand a little bit about a lot. And I know a lot of people that I can go to and ask questions of. But I'm not the average person in the business community at ICANN. Gee, I was expecting laughs at that. Thank you!

Your reports are rich with information, but there's a very short list of people who come from the common everyday need to be engaged in making policy at ICANN from some of the sectors, like the Business User Community and the Civil Society etc., who can easily digest the information and apply it and extend its reach.

This is not a criticism of what you do. It is in fact an issue about how perhaps additional resources could be supported by ICANN to provide a different kind of translation. Not just a language translation, but a translation in terms of the understandability about how users, whether



they're in small corporations or large need to think about how to apply the recommendations that are in your reports and the risk that might exist if they don't get engaged.

There's a particular sector that's very actively engaged in the SSAC. Comcast has been a longstanding contributor to the SSAC, but there are many, many countries where a sector is not yet up to speed on how to get engaged and/or why to get engaged.

So I think it's not just the business users around the world who could use this information. The more we simplify the ability to connect to the Internet, the more we introduce risk. I once had a conversation with an official from possibly a country that will have the largest or the second largest population of end users, and he told me how much he appreciated a quote that my previous CEO had made, that when he was asked what the greatest risk was to the Internet, he said, "It's the backhoe." This official said the greatest risk to the Internet might be the uneducated user.

So I think we want to make users not just users, but informed enough not to be at risk or put other people at risk, and that takes a different packaging of your information.

PATRIK FALSTROM:

Thank you, Marilyn. I'm actually very, very happy to tell you that we actually do have an answer to this very long question because we're already working on it. So let me ask Jim Galvin, our vice chair, to explain what we are doing.



JIM GALVIN:

Thank you, Patrik, and thank you Marilyn. Something which we have just started is working with ICANN's communications staff. Our admin committee, we have regular face-to-face meetings in addition to the full SSAC committee, and one of the things that we engaged in during our last meeting with staff was a discussion of how to better communicate out SSAC's work products and the various activities of SSAC.

So, in fact, we're going to be developing additional summaries, additional translations; and, in fact, targeted presentations of our work and our activities. We'll be creating some video reports, little webinars, and recording those and making those generally available. So the details of this are actually just a work in progress, but in fact, just to expressly more generally answer your question, we are looking for better opportunities to do some outreach and make it visible, exactly what SSAC does and how we do it. So you'll begin to see that. I don't know if you'll see anything by London, but I would expect by the end of this year you'll see some visibility into those activities. So, thank you.

PATRIK FALSTROM:

Let me add that I got some specific questions also when I met with fellows and Janice that we are so early in this process, so we are planning how to produce this material. But we are still thinking about how to do the distribution, so I'm happy to – Steve Sheng is the staff support on this. Happy to talk to you, Marilyn and others, on how to try to do the distribution in as effective way as possible.

Because what we have seen when we look at our reports, the reports that have been the most effective ones are the ones where we choose a



language which is targeted at the kind of organizations which there are recommendations targeted at.

So if it's a technical report with technical recommendations, it's good if the language is technical. And vice-versa. If it's towards lawyers and governments, the language should be towards those.

This is also another argument why we feel we need to do communication that targets the other groups that the recommendations themselves are not directly targeted at.

JIM GALVIN:

Just to give a little bit of credit here, special credit to our support staff. It's important to point out. It was Steve Sheng, one of the three support staff we have, who had started this whole project. He had actually brought the suggestion to us to begin to do greater outreach communications, in part because ICANN is doing more of this. So he gave us that opportunity and we are jumping on it, so I just want to make sure to thank Steve for taking point on this and making sure that all this happens.

PATRIK FALSTROM:

Any other questions? Marilyn?

MARILYN CADE:

Among other things, I am one of the officers of the Business Constituency, so I'd like to volunteer to be a guinea pig in any way that we can do to help to provide feedback or outreach. We have a number of associations that have very wide reach geographically and we'd love



to do whatever we can in helping to distribute the information, but also in giving feedback about how it's been received if that's helpful.

PATRIK FALSTROM: I'm looking at Steve. He's nodded, so that is noted. Thank you very much for the offer.

UNIDENTIFIED FEMALE: This is a comment or request from Tyrek in the chat room. He says: "The SSAC reports are very useful, but in my country, Bangladesh, not enough decision-makers speak English. Please consider working with the regional offices of ICANN to distribute and work on your reports. This is not best done by a centralized function, which is what you are suggesting."

JIM GALVIN: Thank you for that. I'll add that one of the other things that we are doing is we will be looking specifically at translation and better distribution into target markets. We don't expect to fully translate all SSAC reports. To that end, what I would ask is if you have a special need or a special desire to have access to a report and you would like it translated, that we would certainly like to hear that information from any community or audience that would like to see it.

Our plan at the moment is to translate executive summaries of reports and make those generally available, and we will from time to time target specific reports for broader translation, but for the serving the needs of the community, if you have a special request, please do let us



know and we'll see what we can do about getting those reports translated. Thank you.

PATRIK FALSTROM:

Let me ask people in the room, how many of you are using the interpretation that we are using? If anyone is using interpretation, please stand up. Thank you. I do know that we are a [technical] community which normally uses English as the working language, but I still would like to have a straw poll, because for us doing technical work, I am looking [inaudible] because there is a budget constraint. Interpretation costs a lot of money. Translation costs a lot of money. So I'm looking at of course the balance for us in SSAC on maybe a balance between the two. I know it's very, very sensitive to mention it, but at least I think it's important to gather some data.

The feedback we have got so far is that translation of what we're doing is more important than interpretation, but we're still gathering data, so don't draw any conclusions of what I just did.

Bruce, please?

[BRUCE TONKIN]:

Yeah, Patrik. One thing in terms of which documents to translate, I think if I look at the SSAC reports, some of them are more informational or tutorial in nature, such as some of the reports on DNS hijacking, and other reports are really aimed at operators to make some changes to their operations. So you might consider perhaps slightly different audiences, a little bit like what Marilyn was referring to, creating some documents more for policy making or senior executive audience and



other documents that are more geared towards the actual system administrators that need to actually make changes and configurations and things.

PATRIK FALSTROM:

Thank you. That's good feedback. As [inaudible] said, we're also thinking about maybe being able to translate at least the executive summary or many more reports as well, and then also coming to a process where we listen to the community. So if we publish a report in English and we get strong feedback for one of them – "Oh, this is very, very important" – then we can translate. English is still the main language and where the authoritative data is, and then we do translation when we see that being needed. So, yes.

So feedback is extremely important for us to know what we're going to translate, so please let us know if there is something you want us to translate. Three more minutes. I think there is one more question, Julie.

[JULIE]:

Thank you, Patrik. This is a question from Ralph in the chat room. He says: "Concerning the advice to network operators in Merike's report, isn't this something that the GNSO Council can adopt and work on with their ISP constituency? There should be tight coordination so that GNSO can take action and things don't wait just for the Board."



PATRIK FALSTROM:

I can happily report that the GNSO has started an activity in the more close orderly way and ordered way keep track of the recommendations from SSAC and see that they are acting on it.

The Board has already started to have a tracker for recommendations from SSAC and ALAC, and what we now see is that we have very good communication with GNSO where they have a mailing list or work group. I'm really nervous about the terminology here, but they start initiative with the work group which have a mailing list which also, for example, me and other SSAC members are on where we can explain and participate in a discussion on how they can pick our recommendations.

So thank you for the feedback. That's very explicit. I think we are moving in a positive way regarding the cross-community adoption of each other, SO and AC's, recommendations. More can be done, but I am very happy to see that things are moving forward now.

So, with that, I close this meeting. Thank you very much, everyone, for coming and safe travels home, everyone.

[END OF TRANSCRIPTION]

