COPENHAGEN – SSAC Public Meeting
Wednesday, March 15, 2017 – 15:15 to 16:15 CET
ICANN58 | Copenhagen, Denmark

UNIDENTIFIED MALE: For the record, it is Wednesday, March 15. We are in Hall A1 and we will be beginning SSAC Public Meeting at 3:15 p.m.

PATRIK FÄLTSTRÖM: Can SSAC members come to the stage, please? Find a chair that is not taken.

Welcome, everyone. It's a little bit more than quarter past three in the afternoon. Believe it or not, it's 8:00 a.m. Thursday morning. Even though we also have the morning coffee here.

Anyways, my name is Patrik Faltstrom. I'm Chair of the Security and Stability Advisory Committee. Well, the joke I made was just because normally SSAC meets 8:00 a.m. Thursday morning in a room which is much smaller than this so this is quite a change in multiple ways.

I'm surrounded by SSAC members that didn't have any other duties. We see us as one of the most important duties SSAC members have when being at an ICANN meeting is to participate in other sessions and speaking and be active so not everyone had the ability to be on stage. But I was thinking of going along the table to start to my far left and have people just introduce themselves.

ROD RASMUSSEN: Rod Rasmussen, unaffiliated.

TARA WHALEN: Tara Whalen, Privacy Engineer at Google.

PATRICK JONES: Patrick Jones.

JAAP AKKERHUIS: Jaap Akkerhuis, NLnet Labs.

ROY ARENDS: Roy Arends.

GEOFF HUSTON: Geoff Huston, Chief Scientist at APNIC.

JAMES GALVIN: Jim Galvin, SSAC Vice-Chair and also Afilias.

PATRIK FÄLTSTRÖM: Patrik Faltstrom, Netnod, SSAC Chair.

RAM MOHAN: Ram Mohan, Afilias and SSAC Liaison to the Board.

BEN BUTLER: Ben Butler, GoDaddy, Jack of All Trades, Master of None.

WARREN KUMARI: Warren Kumari, Google.

JOHN LEVINE: John Levine, loosely affiliated with the Internet Society.

JEFFREY BEDSER: Jeff Bedser, iThreat.

GREG AARON: Greg Aaron, iThreat Cyber Group.

ROBERT GUERRA: Robert Guerra, Privaterra.

JULIE HAMMER: Julie Hammer, unaffiliated.

DANNY MCPHERSON: Danny McPherson, Chief Security Officer at Verisign.

PAUL EBERSMAN: Paul Ebersman, DNS and [inaudible] victim at Comcast.

CRISTIAN HESSELMAN: Cristian Hesselman, SIDN .nl registry.

PATRIK FÄLTSTRÖM: Thank you very much. Then we have three out of our four brilliant support staff. Wave a little bit. Hello.

What we're going to do here is to give you a brief overview of SSAC. We're going to look at the work we have in progress, some future milestones, look at the publications since the previous ICANN meeting, and last we will talk a little bit about with you and have Q & A which normally is the most fun part. That's why I'm happy to have so many skilled SSAC members around me.

We are, at the moment, 31 members. The members are appointed by ICANN Board and we advise the ICANN Community and Board on matters relating to the security and integrity of the Internet's naming and address allocation systems.

This is something that we are coming back to in various reports now and then that the scope is SSR-related and related to the identifiers and the naming and address allocation systems.

We have lots of different expertise in various areas and as you see on this slide it's even overflowing. We know lots and lots of things. The important part of this is that when we appoint new members of SSAC, which we do ourselves, we try to ensure as good as we can, to make sure that SSAC as a whole has the expertise needed to be able to write good reports.

We have so far 91 publications and actually, we have the 92nd one that we distributed this week but we have not published it yet. They are divided into reports, advisories, and comments. And all of those together that is what SSAC believe, and think, and say.

Then, we have all of these skilled 31 SSAC members and, of course, all of them and many of you in the audience and many others can speak on security related matters but what SSAC say is what is in these reports.

The Charter of SSAC is bound to the ICANN's Mission and Core Values which is to ensure and the stable and secure operation of the Internet's unique identifier systems and preserving and enhancing the operation stability, reliability, security, and global [interoperability] of the Internet.

Our primary role is to give advice to ICANN Board and that is how what we are doing, our recommendations and our Charter is tied to ICANN's Mission.

When we are submitting a consideration of SSAC advice to the ICANN Board, we submit the advice to ICANN, Board acknowledges and studies the advice, they take formal action, and then they can do one of four things as part of the action.

They can feed the document and advice into the Policy Development Process. They can ask staff to implement the advice with the help of a public consultation process. They can disseminate the advice to the affected parties or they can choose a different solution while

explaining why they chose a different path forward than what SSAC suggested.

So formal action does not mean that the Board has to do what we want and no one has to do and follow what SSAC advice is but, of course, just because it's the best advice in the world and if you don't follow our advice, then, of course, the sky will fall over or other things. Who knows what will happen.

The serious part of this is that we're just an Advisory Committee and all of you should, of course, read our advice and then if you believe the advice is correct, then you will follow it so the advice will stand on its own.

The current work parties we have are related to rate limiting of some data in WHOIS, Harmonization of IDMs. We are following the work and management of the namespace and the risk of delegating new TLDs. That actually did result in publication of SSAC 90 and SSAC 91. It is work that we're looking at whether how much else we should do there and there might be some other kind of work items popping up.

We have ongoing work parties related to the DNSSEC Workshop and that we run – it was today and we also have the membership committee that looks at and evaluate both SSAC members and potential members that apply for SSAC memberships.

Since the last ICANN meeting, we published documents 85, 86, 87, related to the EPSRP which is — sorry, 85, 86, 87 were responses to the gNSO PDP Working Groups; 88 and 89 were responses to the ccNSO

comments on SSAC 84, which is related to the ccNSO EPSRP process. Then, SSAC 91 are comments on identified technology health indicators and SSAC 92, which is on its way out through the door, is related to UMRI issues. So this was something that we already talked about.

If you look at the current and future milestones, as you can see, we have published basically what we were hoping in Q4-2016. We still have a few things that we are looking at during Q1-2017 that we have not delivered on yet and then we are planning for Q2 DNSSAC Workshop at ICANN 59. We're looking at the future work on namespace and new gTLD round.

A little more in detail in the publication since the last ICANN meeting and take them backwards. In SSAC 91, we reviewed the presentation on Identify Technology Health Indicators and provided a response to the call for public comments on the description of the five diseases that could affect the health, the name, part of the system with unique internet identifiers.

We in SSAC have met with the [inaudible] of ICANN a few times and also not only submitted these comments in the form of SSAC 91, we also discussed with the [inaudible] including this week. We have some issues with, first of all, the choice of terminology which now has been updated slightly but also we have some issues with the lack of a clear distinction between collection of the data and drawing conclusions from the data.

In SSAC 90, we have advice on the stability of the domain namespace where we do some observations and recommendations directed and mitigating clear identified risks. In 88 and 89, we clarified some issues that we raised in SSAC 84 related to the EPSRP process within ccNSO related to reevaluation of a failed evaluation of an application in the IDNC ccTLD Fast Track. ccNSO has come up with a report regarding the EPSRP process.

We had some issues with certain issues related to confusability in the report. That is an ongoing discussion between ccNSO and SSAC and after SSAC 89, we have also had several meetings between ccNSO and SSAC also this week including this morning so there's progress regarding moving forward on this topic.

These are examples of questions that we get asked on how we prioritize new work and how we address requests from the ICANN Board and the committee. I was thinking going through them a little bit quickly and then I would call all of you to walk to the microphone and ask whatever questions you want for the remaining 45 minutes.

SSAC prioritized new work and chose what work item we pick up on our own. The difference in priorities is that if it is the case that we get a question from ICANN Board, in that case, we do prioritize that higher than other things because our primary task is to advise the ICANN Board.

It's also the case, of course, that if there are public consultations which has a timer on them that runs out, in that case, we try to deliver

the report in time for the end of that timer and we have been relatively successful of being able to submit comments within the timing period.

We also answer questions and give recommendations based on issues that are sent to us from other constituencies in ICANN. I must, though, say that we have got, last year approximately zero such requests for evaluation of issues. We really would like to have more of those because the last category are issues that we come up with ourselves.

When we investigate and watch all SSAC members see what's going on in the internet and what's going on in ICANN when they are running around here in the rooms, for example, [follow] mailing lists. So we pick up, also, work items on our own.

We also get questions on how we track the Board's response to SSAC advice and specifically advice that we give to ICANN Board is important for us because the Board, as I showed earlier, they are required to take action on our advice.

We have been following that manually, historically, which we know and detect it is not really as good thing and that is something that both we and ICANN Board, we agreed that this is not a very good thing. So ICANN Board is developing something they call Board Advice Tracker that will be part of ICANN normal website.

It is the actual user interface to the tracker. It's not ready yet. It will be ready shortly but the actual machinery behind it is ready which means that we, in SSAC, just like others that have interest of getting reports on how our advice and advice from other Advisory Committees are

where their advice is in the process. In ICANN machinery, they can request a report nowadays because all the tracking is now automatic.

The next question. How SSAC informed the community on its work? We do it, for example, now when we explain to you what we have been doing and what we work with but it's also the case that we release the reports. That's the most important thing, of course, we do have our web page where we publish our documents and we see on the statistics of the page visits on the subset of the ICANN web pages that cover SSAC. We see that the web page we have with our documents is the most interesting one for all of you and all of us that actually are interested in SSAC and that's a good thing.

I think it was the case that everyone that goes to our documents web page spends an average of three minutes, or something, on our web page which I thought was good.

It's also the case that we have a Facebook account. We also try to make some videos. Videos that we are releasing with the help and together with ICANN Communication Team. Of course, we are trying to be better on participating in other meetings where we also try to have various SSAC members present what we're doing. So if you have a conference, or whatever, don't hesitate reaching out to us. We have people living all over the world so if you want to have someone present approximately what we're doing now or at presentation on one of our documents, don't hesitate to reach out to us.

We are the security people and not Comms people so I would like thank Duncan and his team for actually—thank you Duncan—for actually helping us understanding how to communicate.

The last question regarding how the community can track the Board's response to SSAC advice. Well, that is also sort of tied together with what I just mentioned to you. It's not really easy to track that without asking, getting some help from ICANN staff but that's what we have at the moment.

So, over to you. Here's some feedback that we would like to have but with this on the screen, I would like to open the microphones for questions and we're happy to answer any kind of—well, ask any kind of question and then we'll see what we'll say. Please.

Okay. Let me start by having and ask anyone is there any SSAC member that would like to bring up a topic. I think I heard one here to the left.

UNIDENTIFIED MALE: You know, it was jest, but it's actually true. We've seen in the last six to nine months the denial of service attacks now getting to absolutely salvage proportion. We're constantly told the Internet of Things is now turning hostile. This is not the Internet we've dreamt of. It seems a very warped and perverted and very dangerous internet.

What can we do about it?

PATRIK FÄLTSTRÖM: If you, in the audience, go to the microphone, you get a membership of SSAC. Anyone?

So, regarding—oh, good. Excellent.

JONATHAN MAKOWSKI: Jonathan Makowski with RiskIQ. Thanks for the great work you are all doing for the community. Just a couple of things on my mind I wanted to bring up and get feedback on.

One relates to DNS sinkholes and whether or not there should be any guidelines from the committee regarding the use of sinkholes whether it relates to WHOIS inaccuracy or the potential for drops. I know the Swiss government, right now, is working on legislation in that area and I'm wondering if it's been discussed and whether you think it's a topic of importance to the community from your standpoint.

I have other questions but maybe I should—

PATRIK FÄLTSTRÖM: Let's start there. I actually don't know if anyone here are working with—Rod?

ROD RASMUSSEN: I actually talked to you about this earlier, Jonathan. This actually was the—the sinkhole topic was actually the genesis of a failed work party we had several years ago but it may be time to revisit that. Especially—

I was unaware the Swiss government was planning on doing some legislation around this.

For those of you in the audience who are like, "What's a DNS sinkhole?" The concept here is if a domain name or I guess any host name is being used for something like a malware command and control so the bots that are out there on the Internet look up a domain name to get instructions and then carry out those instructions based on what that server tells them to do.

What happens if you are law enforcement or security researcher or a security company, etc. is that those domains will be taken down by the registrar or whoever had the resource at some point and either put back into the pool for people to register again, which sometimes the bad guys just go and do that. But sometimes the good guys go and do that and they—or they may have some arrangement with a registrar and they will take that domain name and point it at what is called a sinkhole.

What happens is all those infected computers was still continuing reaching out to that same server and with that communication, you can tell what computers are out there that are infected in the world. Then, depending on your aims and purposes, you may let ISPs know about that or hosting companies through notification, you may need to use it for research, you may do both. There is lots and lots of examples.

Some companies have built their entire business based on this kind of data and then turning around and selling that for information that

people then around and use for their protection; maybe some nefarious purposes as well. So that's why the question of the ethics of this come around. There was a case a few years ago where there was one security company that was going to registrars and basically stealing sinkholes from other security companies. That's obviously a thorny area.

The topic remains important. Just another example of a sinkhole is when a malware author writes what's called a domain generation algorithm or DGA for short. There's another acronym for you to memorize. Anyway, with a DGA, what happens there and probably the most famous example that you may have heard of is Conficker, which ICANN was very instrumental in dealing with several years ago now.

What happens there is that the algorithm that is sitting on that bot will have a whole list from a few dozen to many thousands of domain names to reach out to, to try and get commands from it. Those domains are not registered, for the most part. So this causes all kinds of issues but it's a really good thing for security researchers to be able to use because they don't have to take over an existing domain that's being used for malware control. They can just register the domains that they want ahead of time, again, sinkhole these so that they get—do the same kinds of things with them.

Policy around that, too, is loose, put it that way. There is an effort that you may have heard of at other ICANN meetings that's been talked about is called the Registrar of Last Resort which is being run by Shadowserver. They are trying to normalize sinkhole operations so

that various security companies, researchers, etc. can use that Registrar of Last Resort as a place to put these toxic domains and then have some rules around using that data information. But that is still really just getting off the ground and actually Avalanche takedown that some of you may have heard about in the news.

That was something that was started in 2008 and finally, they've arrested most of the people behind that, almost 10 years later. Those domains were handled by that system. So it's an area right now where there is some work going on but certainly, if there's interest by the governments to regulate that then it's probably something we would want to take a look at.

DANNY MCPHERSON: Patrik, can I come in as well?

PATRIK FÄLTSTRÖM: Yes, please.

DANNY MCPHERSON: I was going to pile on to what Rod said. I certainly agree with that. We deal a lot of operational activities like the Avalanche takedowns, some of that and others involved in some pretty high-profile activity that impacts lots of consumers on the Internet like Ransomware and certainly every kind of cybercrime or other types of malicious activity you can imagine.

I think that as far as parameters go as long as the namespace was used to help people rendezvous to either navigate on the Internet or to effectuate some kind of malicious activity, then certainly we've got to do what we can in this community for consumer safeguards on that side of that equation.

I think there roller aspect, the Registrar of Last Resort is one of the new ways of looking at that and it still has a little work to be done, I think, as Rod was saying. But it's trying to find ways to be more efficient. At the same time, there's a Public Safety Working Group that a number of folks are involved with is looking at Spec 11 changes for consumer safeguard and so forth, as well.

Then, the other thing, I guess, is there are a few documents have been written out there that talk about how the Internet just works and it's sort of some areas that aren't over prescribed or codified. I think a lot of the information sharing that occurs today and the activities that happen certainly in cooperation with law enforcement, where appropriate and so forth, is people trying to have some long-term more strategic impact. Right?

One of the concerns that exist, of course, is my day job where I have to protect a lot of operational infrastructure, I keep buying equipment like firewalls and middlebox and other things that have growing lists of technical identifiers, domain names, and IP addresses and file signatures and that kind of thing. Those aren't getting smaller and a lot of those things may become scorched earth and never usable again.

So the hygiene of that namespace and long-term effects that allow us to clean up that is certainly important from an operational perspective for the infrastructure as well. I guess SSAC has done a number of broad things around this and some of it is trying to help with the various ICANN folks about the contractual capabilities they have. Then, from an operational perspective either on the Internet security side or registry operations and so forth figuring how to best combat some of these threats that lots of folks face.

If there is something that SSAC should say or do specifically related to this week it certainly revisited but as Rod said we did have a Work Party and talk about this and some of it's so vast right now and it's only been getting more and more problematic with global namespace. So I think that there are certainly some challenges between that and privacy as well. Anyway, I just wanted to help.

PATRIK FÄLTSTRÖM: Thank you, Danny. Anyone else. Unfortunately, the number one SSAC member that is working specifically with the registry—that fallback registry that Rod was talking about, Benedict, is unfortunately not here. Otherwise, we could probably hear much more from him.

Please, next.

DAN YORK: I don't have a question I was just going to more say you mentioned in your list of ongoing activities the DNSSAC Workshops; I just thought I'd give you a quick update that did just conclude. A number of people

here were there but for those who were, there were about probably 100 people there over the scope of the day and a number of good presentations that I'd encourage SSAC members to go and take a look at that were there.

We had a panel about deployment of DNSSAC within European. We .de, .at, .dk, and .cz were in there as well.

We also had a session about ISPs. Paul Ebersman was there representing Comcast, talking about what ISPs are doing to be prepared for the KSK rollover and some pieces like that.

Then, we had a couple of different demos and presentations around how DNSSEC and DANE are being used for e-mail, for securing e-mail and I think one of the cool parts that comes out of SSAC helping support this and promote this was that we had people in the room from different projects who were then getting together to talk about how they could make those projects better together.

Also, I would say for those who are interested in statistics and stuff, Roland van Rijswijk from SURFnet, did a presentation about ECDSA deployment and Elliptic Curve cryptography basically within DNSSEC and pieces and it's loaded with some good stats and info for people.

Jeff, your stats were called out a number of times and Roland is also interested in coordinating an effort around doing some measurements of DNS validation as we approach the July 11th and the new key being in there and then also again, of course, around that.

The recording is there, the slides are there but I would like to just say thanks to SSAC for their continued support of that and it was a great session today.

PATRIK FÄLTSTRÖM: Thank you very much, Dan, for helping and let's just say that we have absolute no intention ending this cooperation. The contrary, let's continue to do the right thing here.

Next, please.

MARIA HALL: Thank you. My name is Maria Hall. I'm here as a RIPE NCC executive Board member but I'm also Chair of the Swedish SSAC chapter. As Jeff was talking about Internet of Things, I couldn't help myself to want to have some, maybe you can elaborate a little bit of this because I want to know—sorry, what is your name again? Yes.

ROD RASMUSSEN: Rod Rasmussen.

MARIA HALL: Okay, Rod. I would like you to elaborate, you or any other of you of SSAC board, elaborate a little bit of what the Internet of Things development as you kept caught up [inaudible] has to do with—it's connected to actually to what you were elaborating on. And, of course, it is connected but I would like to think or hear a little bit about your

ideas about that. As Chair of the Swedish SSAC Chapter, we had an Internet of Things session actually last week. We had the Board meeting and we had the General Assembly. There was a guy talking about some kind of—I don't know if it was a regulation or legislation, some kind of technical requirements for all the vendors. I mean, either toothbrushes or they are [fridges] or there are a bunch of things connected—not to talk about our own laptop and everything. We were talking about that a little bit but I don't know if you heard about that.

PATRIK FÄLTSTRÖM: One thing that we have been looking at while the rest of you prepare, one thing that we have been looking at in SSAC a little bit but it has not resulted in the report because we don't really know who to give the recommendation is that one of the largest problems with all of these Internet of Things is that people buy them, they connect them, and then they don't touch them and many of them—which means that people don't upgrade the software even if bugs are found. There are too many things where you cannot even upgrade the software even though you would like to.

So, that is another area regarding legislation, to enforce people, it's like telling all the criminals that they are not allowed to break into houses, and so far that has not really worked that well. Is there anyone who would like to add on? Danny?

DANNY MCPHERSON: Yes, I think I would just say there are a couple of places where certainly some of the big DDOS – I think Geoff's question was fair, by the way. I think that we probably should say something about this at the public forum panel.

So, I think that the registry operators that are involved in various places in the ecosystem, if you're at the top-level domain or the root infrastructure, or somewhere at a second level domain, like a [inaudible] that's on the receiving end of a lot of packet love from an attacker, I guess, but large scale attacks, we obviously take this stuff very seriously and continue building infrastructure and looking at new ways to collaborate and preserve the availability and integrity of that infrastructure.

When network and navigation and trust is your business and you see an attack like that that has the kind of collateral damage that it does, it's really problematic. So certainly, the partnerships that the ICANN community helps bring together and also the operational security community is one side of the spectrum.

Another is some of the prior work all the way back to like SSAC 04 with anti-spoofing which is one of the attack vectors in some of the IoT-based attacks that we saw right there. There are a number of SSAC documents that sort of touched the entire landscape of some of the DDoS activity related to this.

And then recursive nameservers are an example, where some of those were embedded or encoded in some of the attack code that the Mirai attackers in particular used. So, that's another side of it.

And then the broader operational security community. I'm not sure how much of this is actually an ICANN thing per se, although I think a lot of folks that participate in this community have some influence there. But there are some IoT trust frameworks and other things that a number of folks like the Online Trust Alliance and even FTC and NTIA and others are involved with related to these kinds of activities, at least in North America, and some of those are more global as well that talk about best practices.

But I think that there are more devices out there and more things we care about that live in more places with less capability to protect, then we certainly have to try and continue to raise the bar and make the upgradability and security of those devices paramount, and at least controls to mitigate any attacks associated with that kind of activity, in particular when they intersect with the namespace and the number space as well on the Internet.

So, I can say that we invest extensively in capacity to be able to absorb DDoS attacks and we also look at partnerships and other things that would provide extra capabilities because of just that, and so we obviously take it really seriously.

PATRIK FÄLTSTRÖM: Geoff.

GEOFF HUSTON: There are two parts to this. One part is the software and hardware industry, which is in a complete market failure. There is no market for

high quality software out there. There is a market for cheap stuff in the Internet of Things, and that cheap stuff is unfixably bad.

Because we have global supply channels, any kind of regulation of that supply channel is completely ephemeral. We are stuck with a situation where the most pervasive botnet around, Mirai, was actually recruiting people on Telnet, a protocol that one one's used for 30 years, or so we thought, apart from video cameras.

This is unfixably bad, and there is no mechanism, considering these global supply channels, to even come close. What we have observed, however, is there are only two protocols that work on today's Internet. Only two: HTTPS and the DNS, and that's why the DNS is kind of the point of vulnerability.

There's no point setting up a botnet unless I can rent it out b the hour, and the command and control channel is basically the DNS. It's real time, it can actually script literally millions of devices to do my bidding, and the DNS is always there.

That's where if you do all this work on, can we make the DNS signal that there are such things around and even interfere with those command and control channels? It's the work of desperate people. Desperate because the supply side is overwhelmingly big.

Seven billion devices out there on today's Internet, and it's simply getting worse every day, let alone every year or every decade. So, what can we do? We can hope that we can do something productive in the

DNS to try and identify and understand these control channels. We're trying. Are we winning? No, but we're trying.

PATRIK FÄLTSTRÖM: Rod.

ROD RASMUSSEN: Just to not be quite so bleak as my good colleague Geoff is, I think there are some hopes here. But just to talk about the fundamental problem here – and I think Geoff touched on it really well – when you think of the Internet of Things, it really is just more stuff on the Internet than we've had before. But it's a couple orders of magnitude bigger, so we've got a couple of problems to deal with.

One is just scale. We're growing the Internet by a couple orders of magnitude over the next ten years. I think that's one issue, and then another issue is the misuse of that infrastructure. And the issue with those devices, part of it is the manufacturers and the software companies. These are not actually software companies at all. They're grabbing libraries and things like that and jamming them in there, and they're actually including things they probably shouldn't as part of that technology stack, and they're not using good practices.

For all the years that we've been bashing on Microsoft for driving insecure code, they're really doing a good job these days in general, but they're not the ones building these devices.

There's a lot of work. Danny mentioned a few of the places. There's a lot of normalization work going on to take a look at that, and I think there are things that could be done with supply chain. We made toasters safe – before we put them on the Internet, that is. So there are analogies in the real world around that, but I also think that to take a look at this, we're not going to solve the device level problem, and it's a huge scale.

We've got way more devices than we have human beings on the planet. So, that's a harder one to solve, but there are control points. There are access points. If they're going to be on the Internet, there are things we can do to harden that infrastructure, where we can look for things that are connecting to the Internet in ways they shouldn't.

Should your toaster be talking to somewhere in Kazakhstan? Things like that, that nobody is doing right now or very few people are doing right now. So, there are a lot of small lifts we can do as operators of infrastructure, as people who control networks and things like that.

So, there is some hope to at least calm this down a little bit. I think Warren really wants to say something over there.

WARREN KUMARI: Warren is desperately trying to decide if he does want to say something or not. One of the problems is also the term, the Internet of Things. It means very much whatever you want it to mean this week, and it changes next week.

But as a bunch of people have said, it's a bunch more devices, and it's a bunch more devices made as cheaply as possible because in general they're trying to roll out many thousands of them.

Often, this is being made by small companies who don't have software as their main sort of technology or competency set, and they just sort of cobble some things together and hope that it flies.

Writing documentation on things like BCPs I don't think is going to accomplish anything. These folk aren't really trying to follow best practices, they're just trying to make a light bulb that works.

I think also trying to have any sort of control in a supply chain or trying to profile the device to see where it's talking to is also not going to work, because if you look at many of these devices, they connect to lots of things and it's really hard to figure out why.

They have ties into various other things, like If This Then That or various Amazon AWS services, profiling and to understand why they do what they do is really tricky.

What I think might make things better is one of the things which has sort of improved some of the CPE, and that's people no longer – or in general, manufacturers don't write their own code for CPE. They take one of the existing pieces of software, like Tomato or any of the other things.

UNIDENTIFIED MALE: What's CPE, Warren?

WARREN KUMARI: Sorry, customer premises equipment, like the little home router thing that you buy. Your firewall, your gateway. In general, people aren't writing the code from scratch for that. they're taking existing open source code and kind of changing the front page so that it says Linksys instead of saying NETGEAR or similar.

If we could create a framework or toolset that allows people to easily build an IoT device, a thing where you want to make an IoT device, you download this piece of software, select the modules that you want, put your logo here, and now it will look like your thing. That will make it cheaper than people building stuff from scratch.

So, if we can create a set of tools, chains, frameworks, etc. that manufacturers can use and have an incentive to use because it's cheaper than building it from scratch, we might be able to make some sort of a change to the whole ecosystem. Maybe.

PATRIK FÄLTSTRÖM: Next, please.

UNIDENTIFIED MALE: Thank you. I want to broaden the discussion a bit, because on the other end of this, while 25 years ago I founded a security company and in those days we would have patched a fridge and would have patched a toaster, but today, we have the experience that some

devices which are very important are not allowed to be patched because of compliance.

So, compliance is on the one hand a fine thing that you prove that you are secure, but then you are not allowed to patch the x-ray in the hospital or the whatever running machine in the power plant because of compliance reasons.

This is the other end of the Internet of Things, because it's connected today and it's unsafe, you're not allowed to patch. So, what is your comment on that, and how do you deal with that?

That's the first, and the second is, you talked about the security of HTTPS. Yes, but our customers want interception, they want to look into that, so we get NG firewalls which just split up everything, make man in the middle attacks, and so all the security is gone, our users see a green light and think it's green, but it's not green because they're intercepted. So, how [to deal] with that? Thank you.

PATRIK FÄLTSTRÖM: Regarding your first issue with licensing and compliance, that is as you say yet another piece in the puzzle. And this is one of the reasons I think as the Chair of SSAC why for example we in SSAC have not really found sort of where should we start.

The problem place is so large, so it's literally very hard to focus, and also to try to give some kind of recommendation that makes the world better like Warren was talking about. Because writing documents

about the problems with Internet of Things as we see here is pretty darn easy.

We in this room could probably write a 50-page article in just the next hour, but from that, conclude what we think has the best effect is really hard.

Regarding the second thing that you talked about, the SSL or TLS connections, we in SSAC have issued a number of recommendations regarding how to make and use certificates in a more effective way, weaknesses with the certificate, CA mechanisms that we're using today.

We have responded to ITU-D that sent us a liaison where they asked various parties whether they thought it would be a good thing to launch new CAs all around the world so there would be a higher number of CAs, and our response was that, no, it's actually the other way around.

We are referring to also DANE that we heard mentioned earlier. So, I think the SSAC view that TLS connection, that encryption tunnels like TLS and similar mechanisms should continue to be end-to-end, and for example use DANE and similar technologies that we have said that several times.

May be a little bit hard to find in all our documents, but you can see that as a common thread whenever we talk about certificates.

Anyone else? Please, come to the microphone.

UNIDENTIFIED MALE: [inaudible] University of Oxford. I would be interested, what would you be interested in researchers researching from your background and security? What are the things we should address?

PATRIK FÄLTSTRÖM: Is that something for you, Geoff?

GEOFF HUSTON: I can certainly respond. I too do a lot of research, and I must admit, my particular pet project is to understand the Internet from the edge looking in, not from the in looking out.

It's always easy to measure your own infrastructure, measure your own server, measure your own something. And it's easy to set up a little laboratory, but it's hard for me to be you and the other three billion or so like you, and look at the Internet the way you see it.

But we've been successful in some areas, and in particular the massive measurement projects we've had going, look at the penetration of IPv6 across the entire Internet almost in real time. And similarly in doing some introspection in the domain name system, looking at how many of the world's users will not go to a domain name if it is badly signed with DNSSEC.

That's a big statement, because it says they have to be DNSSEC-aware, and they have to tell the difference between the truth and a lie. Now, Sweden is brilliant. Huge amount of DNSSEC validation.

My own country, Australia, I seem to be the only one doing it. So, it varies a lot country by country and provider by provider. But that kind of thing about, what does the Internet look from everyone on the edge inwards is kind of the big questions for a lot of this, because the Internet is all of us. It's not the bits in the middle, it really is the people at the edge.

PATRIK FÄLTSTRÖM: Warren?

WARREN KUMARI: And I think kind of following on from Jeff's thing, how the Internet looks different in different countries. Understanding better what things you can reach and what things you can't, which is often a function of censorship or something similar.

Why can't I reach this specific set of sites from a specific country? What is it that they do or don't want me to see?

PATRIK FÄLTSTRÖM: Danny.

DANNY MCPHERSON: One other place I think I've certainly seen a lot of interesting research from our folks who have been involved and I know Geoff and Warren and others who have done some is systemic dependencies that relate

directly to things like your transitive trust graph in the DNS or cloud infrastructure where you may have things that are resident.

I was just trying to pull it up, I don't have it. I think it's around for example of the nearly 300 ccTLDs, about 60 of those I believe was the last number I saw, don't have a single in-country authoritative nameserver for their ccTLD.

So, when you start to look at that and say, "What am I dependent upon for operations in cyberspace and your nameservers aren't even in your country," for example.

That's a sort of simple thing. Or understanding sort of multi-tenant infrastructure like we've seen in Dyn, AWS or other outages where somebody has some impact, and then all of a sudden it breaks 1000 companies, services or capabilities.

I'm not sure a lot of people even enumerate that, it's just going to get worse as more things move to cloud or broader infrastructure because of the efficiencies and gains in that, so I think that both the security implications of transitive trust even in the DNS alone, but certainly for broader Internet infrastructure, and then certainly the national security, state interest kind of thing, and then also business resilience and business continuity is really important, and I think how distributed infrastructure and systemic interdependencies impact that are really important things to look at.

PATRIK FÄLTSTRÖM: Thank you. Cristian and then we'll go to the microphone.

CRISTIAN HESSELMAN: Yes, two more research topics – or perhaps one I think is for end users to be able to give them more control on the security and privacy aspects of the devices they're using in the home for instance, so that is basically related to the IoT discussion we just had, because a potential hypothesis could be that if they were aware of security vulnerabilities, they might switch that device, take the device off the net, and that both protects themselves as well as operators in the DNS infrastructure.

And also, perhaps another topic would be how different DNS operators and other infrastructure parties could collaboratively analyze data and also share that data with each other when it comes to security.

PATRIK FÄLTSTRÖM: Next, please.

WES HARDAKER: Wes Hardaker, USC. That last round of discussion, we did what you did, but what we all normally do – which is we group things into classifications that are easy to identify and measure. Countries is sort of one example.

One really interesting study that SSAC could dive into is how many other measurable categories are out there that we're not looking at.

An obvious one might be upstream [ASes.] What's the percentage of DNSSEC validation at your upstream ISP?

So, countries is one. The reality is that the world doesn't follow country borders, as we talk all the time, and we know that the Internet doesn't, but yet that's still our standard of measurement in terms of binning people.What else might we have available?

So brainstorming into that session. I don't know if I dived all the way down into elementary school and universities. I don't know, but it seems to me like that there's probably a lot of very other interesting metrics that would come out if we could bin people in other ways than just countries.

PATRIK FÄLTSTRÖM: Geoff.

GEOFF HUSTON: You'll find a lot of this stuff when you look from the edge in, whereas a cursory look at the routing table gets you straight into their home network. So, we certainly look at autonomous numbers as being the other key identifier.

Oddly enough, what we found when we publish stats is that if you publish stuff by country, many more people read your publications than when you publish then by this ephemeral thing called autonomous system numbers. So, in some ways, the way we produce

our stats and output is geared towards what we think the readers want to see.

PATRIK FÄLTSTRÖM: Next, please.

NICK SHOREY: Hi, my name is Nick Shorey. I'm from the UK government. We've recently started to look into IoT and this notion of secure by default. You're welcome to make your assertions on the feasibility of such a thing. But one of the things that I've been baking everyone's cookie over was what you do when we get IoT [inaudible], one of the things I was worried about is device obsoletion.

And as you said, you've got all these sort of extra companies whose primary background isn't sort of ICT, and then them creating all these IoT devices. It's all about, as you say, that time to market, push the next thing and fashion. I'm worried about the increase in rate of obsolete devices that are connected, that aren't getting patched and stuff.

I'd really welcome your thoughts and your feedback, either not now but on good ideas to get around that problem where you've got more things online that are obsolete and are past their service life. And if you could point me in the direction of any documentation or research that I might be able to share with the folks back in my government to pick up on this point.

PATRIK FÄLTSTRÖM: Warren?

WARREN KUMARI: I don’t think I've got a good answer for you. I think I've got a depressing answer. I recently was looking for a new access point for my house, and I started looking to try and find a cheap access point. I went on to Alibaba – which is a large online marketplace – and I found about 7-800 different companies or manufacturers selling from what I can see exactly the same device.

It appears you make a couple of hundred of them, and then you fold up shop and make a new couple of hundred of them. So, anything where we sort of expect a manufacturer to stick around for a long time and continue to provide updates doesn’t really make sense for them. It's cheaper just to start a new company.

One possible option seems to be what I think Linksys did many years ago, which is you would buy a home router and would work for around two years, and then mysteriously, the power supply would fail. And then you'd go off and buy another one.

So, potentially when things get cheap enough, they're just going to die from old age and you'll have to throw them out, and maybe the new one will be less bad. Maybe.

PATRIK FÄLTSTRÖM: Rod.

ROD RASMUSSEN: Two things. One, it's even worse than you asked about there, because what do you do about somebody who sells their smart car or smart device, sells you a home that's got a NEST or what have you. How do you change those passwords? Is there an owner's manual for all your electronics in your house?

That's the bad news. The good news is there is a lot of work going on. I would point you right to the Online Trust Alliance which Danny mentioned earlier, because they're gathering up all this information and commenting on it on many of the forums around the world that are actually dealing with this. And there are frameworks out there for addressing all of these issues. These are not new things. People have been really digging in, trying to say, "Here's a standard of conduct for manufacturers, here are regulatory ideas."

I know, Warren, you're laughing at that, but it does work occasionally. See, we don't always all agree on everything in SSAC. But anyways, there are some serious efforts going on there, and I'd be happy to talk to you offline about that.

PATRIK FÄLTSTRÖM: Robert, and then the microphone.

ROBERT GUERRA: Just to follow up quickly on Warren's point, Ondrej who's also from the SSAC who's with NIC.CZ, I'd suggest looking up turris.cz. It's a

Czech ccTLD put together an open source auto updating router which they're slowly rolling out. It's an interesting initiative that also takes the data if you're interested so you can share.

So, there are some initiatives and others are doing that as well too, so it's being recognized by some. So, I just wanted to note that. Thank you.

PATRIK FÄLTSTRÖM: And then microphone. Last question for today.

JAD EL CHAM: Hi, my name is Jad El Cham and I'm a first-time ICANN Fellow. First of all, thank you for the presentations and for your answers. I have a question which I've been asking for the last three days around here while I never got any real answer.

We keep hearing about the Internet of Things and the DDoS attacks and so on, but it seems that we forget that the new types of DDoS attacks are based mainly on two protocols, which are the DNS and NTP where we use the servers as reflectors.

So now, instead of talking about DDoS, we hear a lot about reflective DDoS. I see one of your colleagues nodding his head, but this is what we're seeing in our market, and we do deploy security appliances for our customers.

So my question for you is the following: DNSSEC really does address a lot of security loopholes in the DNS, but I was wondering what is

ICANN doing or advocating in terms of awareness to this shortcoming of the DNS like we can use it as reflectors.

I know one of the answers I got is that there's a taskforce in the IETF which is handling this, but I would also like to know if ICANN is involved in these types of efforts.

PATRIK FÄLTSTRÖM: ICANN itself is just like many organizations, just like [I'll say] Netnod operating some IT infrastructure. They're running one of the root servers just like we at Netnod do. Google is running a lot of the other infrastructure. And all of us as organizations participate as you mentioned for example in the IETF and other operational communities to come up with the best practices on how you configure and how you operate these protocols where there are either weaknesses or that are the most popular protocol of the day to try to use as an attack.

Personally, I've seen many changes over what kind of vectors the Denial of Service attacks are using, but in general, I can say that they have been using spoofed source IP addresses and they're using non-spoofed source IP addresses, all depending on how popular it is and how the various botnets are configured.

Some of the botnets are so large and so hard to find and move around so fast, you simply don’t even have to spoof the source IP address. The last couple of attacks in Sweden against media for example is just normal HTTP traffic. It's not at all NTP or any other popular UDP-based traffic

It's also the case that there are some reflection attacks which are using those protocols just because you also get amplification. Because of that, there are best current practices on for example how to configure your NTP servers so that some of the control commands are not recognized by the NTP server, etc.

So, there are various of these kind of things that happen. ICANN as the organization that operates these things do of course participate just like the rest of us. Regarding the Policy Development Process, there are various different mechanist both on the non-contracted and the contracted party side where these kind of things are taken into account when looking at the various requirements for both the contracted parties and also best practices that they're talking about in all different kinds of constituencies here in the ICANN community.

But regarding operational issues, I think the answer is that we who are participating in the ICANN community and ICANN ecosystem, we also participate in the forums when these operational issues are discussed. Anyone who would like to add something? Warren?

WARREN KUMARI: Yes, I guess the things you're speaking of specifically are reflection attacks, and those almost all require the ability to spoof source addresses. There's an IETF document, BCP 38, which says you shouldn't allow people to send packets that are not sourced from the network that you run. And SSAC published, I think it was SSAC 004, basically saying the same thing. There's not a huge amount that ICANN

can do about it, but who runs the MANRS program? I think that's an ISOC –

UNIDENTIFIED MALE: [inaudible]

WARREN KUMARI: Yes, I just couldn't remember who was helping run it and fund it. ISOC runs something called MANRS which is a sort of opt-in thing that networks can say that they have good manners, that they followed the guidelines in this. And that also says you have to do things like BCP 38, etc.

Unfortunately, ICANN itself can't do very much about this. ICANN doesn't have any real hammer to bang up ISPs with, unless maybe the ISP constituency or sort of the ASO could say, "We're not going to give addresses out to ISPs who don't do spoofing protection," but honestly, that's not likely to accomplish much.

This has been a problem that's been around for a long time and just seems to kind of hang around.

JAD EL CHAM: Alright, thank you.

PATRIK FÄLTSTRÖM: Thank you very much. With that, it's 4:15, [indeed] two more minutes even. Thank you very much for coming, and it seems to be the case

that 3:15 PM on the Wednesday is more popular than 8:00 AM Thursday. Thank you. That's good feedback for us.

[END OF TRANSCRIPTION]