NATHALIE PEREGRINE: The call is now being recorded and I'll do the roll call. Good morning, good afternoon, and good evening, everybody, and welcome to the At-Large Technology Taskforce call on the 16th of November, 2015.

On the call today we have Alexis Anteliz, Harold Arcos, Lutz Donnerhacke, Klaus Stoll, Vernatius Ezeama, Carlos Quintana, Gordon Chillcott, Stuart Clark, Glenn McKnight, Maureen Hilyard, Dev Anand Teelucksingh, Alfredo Calderon, Carlos Watson.

Our guest speakers on the [core] projects are David Goulet. He is on the line with us. Welcome, David.

We have received apologies from Olivier Crepin-Leblond.

From staff, we have Terri Agnew and myself, Nathalie Peregrine.

I would like to remind you all to please state your names before speaking for transcription purposes. Thank you ever so much, and over to you, Dev.

DEV ANAND TEELUCKSINGH: Thank you very much, Nathalie. Thanks for attending this call. I know that some of you are perhaps – this is probably your first Technology Taskforce call. So I thought I'd do a quick overview of what the At-Large Technology Taskforce is about.

So I just put together some slides here. The Technology Taskforce evaluates and reviews ICT Information and Communication Technology that can help the At-Large community – that's ALAC, the RALOs, the At-Large Structures – better able to accomplish their role in ICANN activities.

So anyone interested in ICT and how it can be applied to solve the needs of ICANN At-Large and other ICANN communities are welcome to join the Technology Taskforce. You can just do that by e-mailing ICANN At-Large staff, which is the staff e-mail address.

Just to give some [inaudible] activities done between ICANN 53 and ICANN 54, the two public face-to-face meetings this year, we have looked at Kavi as a possible project policy management process tool. We have looked at technology such as Teamup, which is a group calendar that's in use by the ALAC Outreach and Engagement Subcommittee.

We also had discussions on the LACRALO mailing list issues. We also had a discussion with persons from LACNIC to discuss their policy development process, and look at their tools. We also even had a first look at the ICANN mobile app that was in beta for the ICANN Dublin meeting.

Just to say that it's open to all members of At-Large from all the five regions of ICANN – North America, Latin America, Caribbean, Africa, Asia, Australia, Pacific Islands, and Europe.

Just to quickly conclude, the Technology Taskforce has members from the At-Large community and other members of ACs and SOs. We now

have persons from the GAC joining this working group. We have one or two conference calls per month. There are two links there which talks about where you can find more information. Some of the documentation that we have done for the At-Large community in terms of translation tools and so forth.

I'd like to now go into why we are having this call today. This is one of the At-Large Summit recommendations. For those who don't know, the At-Large Summit was a meeting of all the representatives of At-Large that was held during the ICANN 50th meeting. At that meeting, all the At-Large representatives developed [inaudible] recommendations and observations. It was 43 recommendations in all. You could find it at the link in the presentation.

Some of these recommendations were given to the TTF in coordination with some of the other working groups for implementation. One of the recommendations was recommendation 17, which says ICANN needs to be sensitive to the fact that social media are blocked in certain countries, and in conjunction with technical bodies promote credible alternatives.

So we have looked at group chat services such as Slack and Hipchat. But we also started to look at two that could be used to perhaps circumvent blocked websites.

There's been some discussion within the group as to whether should we consider to use these tools, and we decided that we need to just really have some conference call to really understand some ideas about these tools. And one of the more popular tools is the Tor project. This is

where we have our guest speaker, David Goulet, who is a person involved in the Tor project. He's very kindly agreed to come on the call at this time. So I would like to now turn the floor over to David to give his presentation. David, you have the floor.

DAVID GOULET: Hi. I'm David. I'm from Montreal, so I'm French-Canadian, so sorry about my accent or whatever. You might have some problem hearing me or understanding me, but I'll try to be as concise as I can. I'm guessing you guys have all the slides. I cannot control them, so I'm going to ask someone over the line. I'm going to say next slide.

First of all, thank you very much to Glenn for inviting me. I met him in Montreal and he was kind enough to invite me about Tor project.

We're going to start with the first slide which is the introduction. I work at the Tor project for almost a year now. I've been a volunteer three years prior to my work. All in all, I've been four years with the Tor project. We have a global mission. You can read it there. I'm not going to read it back, but most of it is technology, advocating for privacy. We do a lot of research. We have at least three universities around the world that I know of that do active research on the Tor network, with the Tor network. And we have multiple outreach in terms of freedom of speech, censorship, circumvention and stuff.

If you go to the next slide now, what is Tor? The Tor project created a while back was created by the NRL which is the Navy Research Laboratory in the US. This gentleman, [inaudible], and someone else – I don't remember exactly.

DEV ANAND TEELUCKSINGH: It looks like somebody has… David, let's continue. Let's see if the audio [inaudible] and sorted out.

DAVID GOULET: Basically it was created to have a way of adding a way to communicate over the Internet in an anonymity fashion in manner to preserve privacy for anyone that uses. Of course there was military ideas behind that and now it's moved quite a bit to a project after a few years. I don't know the specifics, but now Tor is a software that's used for online anonymity. It's complete open source. Everything we do is open source. We have nothing that it sells or whatever. All our work is transparent on the Internet meaning mailing lists are [inaudible]. We have open proposals and stuff.

So that community has grown quite a bit in terms of researchers, developers, users, and of course [relay] operators. I'm going to get into what is a [relay] operators.

As of now, our funding is mostly US-based government – so DOD, DOJ. It depends on the year we have grants or whatever. Right now I don't have the specifics again, but most of our money comes from government US. We also have foundations and stuff.

We're working towards [inaudible] source of funding. So if you go to the next slide which is basically [inaudible] for people in the US. It's a 501(c)(3), which is for people in the US – or even the US if you don't know – it's a non-profit organization. In the US, we have an office in

Cambridge and our goal is, again, research and development in the anonymity field.

Let's get into it. If we go to the next slide – slide four – we have this nice statistic here that last time I checked is estimated at 2.1 million daily Tor users. That means on the network we have that amount of users at any time.

If you go to the next slide, you have a graph – slide five – which is directly connecting users. This is only 2015. You can see it's roughly stable at 2.1, maybe 2.2, million users.

By the way, I'm going to use quite a bit of graphs during this presentation and you have the address at the bottom of every graph I think which is metrics.torproject.org. All those graphics are accessible. You can change the dates, download PDF, and so on and so forth.

I'm going to get into the [threat] model of Tor in the next slide. Before I start, there's going to be technical parts and I'm going to some less non-technical parts. I believe we have [inaudible], but Tor is a humungous project, so the goal is really for you guys to understand the whole thing in terms of technology, but also in terms of community and the importance of users, because it reaches outside of the technology part which is the human part.

I'm going to start with this slide which is the [threat] model. A while back – Tor is around ten years old now. We had this idea what an attacker can do when you have this network. In this slide, imagine the anonymity network could be Internet [inaudible] and we have Alice and

Bob connecting. So there's multiple points where an attacker can eavesdrop or either watch or attack one or two participants.

Now we move to the other slide, which is seven. There's a very important thing – important aspect – you need to understand with the Tor network. It is anonymity, not security. By that I mean that we provide anonymity on the transport layer, but then the problem is that you always have the application layer that will [leak] here when you're entering the Tor network or when you leave the Tor network. Fortunately, when you're entering, we fix that but we fix the content leak – but when exiting, that's another problem.

In this case, in that slide, we see that anonymity is encryption. We have "Hi, Bob" "Hi, Alice" blah, blah, blah. It's gibberish. It's still not anonymity because, as we know now, thank you to our friend Snowden is that lots of data on the Internet is being surveilled – [inaudible] basically based on metadata. Anonymity becomes very important here.

If we go to the next slide – slide eight – the thing is that it's not just wishful thinking in terms of technology but also towards the law. If we have government agencies or corporations or government, there's statements saying, "I promise I won't look. I promise we don't spy on our citizens. We're not going to read your e-mail," or whatever. All of this is not enough unfortunately nowadays. We have a bunch of phrases on this slide that just tells you that it's not possible to have anonymity just by saying that it can prove it was me.

This next slide – slide number nine – we have different interests of different user groups. This anonymity cannot just be applied to

government, for instance, where in the first base with Tor and ten years back, it was. The goal was to provide anonymity for either agents on the field or military or just government officials on foreign business.

But the game has changed quite a bit. As Tor nowadays, we have human right activists, businesses, private citizens of course, governments. It got expanded to – you can see [inaudible] in there. Private citizens would be normal citizens. Any one on the field. We know agents from not only the US but other governments are using that everywhere in the world. It needs all those four different aspects here, which is [inaudible], network security, of course privacy. That's the old concept. And traffic [inaudible] is also a nice [inaudible]. And we have those because quite a bit of enemies – adversaries, I would say – that tries to de-anonymize Tor. It's quite a race.

If we go to the next slide – number 10 now. This is the [simplest] design you can see. I'm going to move towards [inaudible] explain why it came to that.

We have Alice and Bob going to Gmail and they relay. Of course this is an issue because if you go to the next slide – number 11 – it turns out that [inaudible], well then, they're eavesdropping the attacks or anything. Censorship, also. [inaudible] single point of failure.

Fortunately, the Internet as we know now, yes, it is a web. It is redundant. We go through different paths, depending on the failures or not. But we still use Gmail in our daily life, which is [inaudible].

Okay, that's the application [rule]. I understand, but still, we count on the very network and so on and so forth that are controlled by single entities. So it's a problem – for anonymity, I mean.

So we go to slide 12 now. Then we have this idea. Okay, we're going to add multiple relays so that no single one can [inaudible] Alice in this case. The more we diversify your path to your destination, we have a reasonable assumption that it's more difficult to attack Alice, eavesdrop on Alice, or do whatever attack is possible. Censorship also [inaudible].

The Tor network – Tor stands for The Onion Router. There's a reason why it's an onion. We're going to see that on the next slide – slide 13 – which is what happens is that each relay… Alice will communicate through each relay and then [inaudible] Bob. So R1 here in this case here would be the entry node, then the R2 will be middle node, and then R3 which is an exit node. And exists [M3] and middle in the Tor network are very different properties. They behave… Well, we hope they behave like the others, like any other nodes, but they still have different roles.

In this case, what's going to happen is Alice will create a path to Bob, say, "Okay, I'm going to pick R1, R2, and R3 and I know those relays from a known document," which we call the consensus. The consensus is created by nine directory authorities in the Tor network. Unfortunately, I don't have a slide for that. But we have right now nine servers around the world that are run by trusted individuals and those servers are responsible for measuring – not measuring, but just creating what we call consensus which is a document that they all agree on which contains all relays in the Tor network that are usable. This

document basically details every relay with their fingerprint keys [inaudible]. That means IP import. And some other useful information.

In that case, when Alice starts up, the first move she does is get this consensus because this consensus is the view of the network. If she already has a whole consensus, she's going to try to get the consensus through a directory cache which is basically any relay in the Tor network depending on some requirements, but mostly all the relays. Or else, if she doesn't have a consensus, this document, she's going to ask those nine directories. And how she knows that is because those directory authorities are coded in the code of Tor.

So the keys, they are public keys and also their addresses and some other useful stuff. It's completely [R-coded]. So that means every client, every Tor relay, everything that uses Tor knows where to get those directory [inaudible].

So back to our slide now, Alice has [zeroed] network. She can then choose R1, R2, and R3. So with that in mind, in the consensus, every relay has a key. A public key is assigned in the consensus. So she knows that she can encrypt a connection to R1 and R3.

If you see in that slide, there's a clear layer. Green, yellow, and blue. The first layer is R1. R1 is the green key. When Alice encrypts, she's going to have this white line which is the content, and then she's going to do three layers for each [inaudible]. Then send [inaudible] back to R1. R1 removes the layer, knows where to send it after, which is information in the packets. Then send it to R2. R2 removes the layer, the yellow layer. And then R3 will remove the last layer.

So what happens here is that R3 is an exit node and R3 has access to the content that was originally encrypted. You see the white line. This is important.

This comes back to anonymity versus security. As a [global server] or as anyone – an attacker being R1, R2, or R3, you don't know where it's coming from. You don't know what is in the data, in the middle. But the exit knows. So the exit knows where it's going and what's the content but they don't know from who it's come from.

So we have this thing on Tor that you really need to use security also in terms of encryption. So SSL, OTR, PGP or whatever you need, so the exit cannot snoop on your data.

If you go to next slide, slide 14, it's a nice onion layer [of the onion layers]. So every relay we peel a layer.

I'm going to leave time for questions, so feel free to ask. It can be complicated sometimes. Okay, let's go to the next slide, slide 15.

This is the number of relays right now we have in the Tor network. I took it last night. That means… I think we have around 6,800 relays. That is 6,800… Well, that means people are running relays – anyone. So a relay is run by anyone that wants to run one. We have volunteers around the world.

If you guys go, please note this down because it's an amazing visualization of the Tor network. You can go to torflow.uncharted.software. If you go there, it's an amazing visualization of all our network around the world.

And bridges here. Bridges, we have 3,000 of those. I didn't talk about those. Bridges are relays that are not advertised into the consensus. The reason is that when you operate censorship in a country or organization or whatever, one of the [inaudible] Tor network – I had a slide about that – is to block all the [relayed] IP because those are known, [because] you can get this consensus.

Now, bridges are not in the consensus, so there's different way to get bridges. I'm not sure if I have a slide about that, but out of my head, for instance, there's a way you send an e-mail to a specific e-mail address and you're going to receive bridges by e-mail. Then you take these bridges, you put it in your Tor configuration and that's going to be your entry point to the Tor network. Since it's not known by anyone nor the consensus, you can evade some censorship. I'm going to come back about bridges with active attackers we have.

If we go to the next slide now, this is the total relay bandwidth, slide 16. You can see right now that our advertise bandwidth, that means every relay advertised bandwidth, we have bandwidth authorities that measure at all times relays and try to make an estimation so we can do much more clever path selection. In this case, we're around 140 gigabits a second possible bandwidth. And this bandwidth [story] is what we think is being used. Again, this is an estimate. This is not that accurate, but it gives us an idea of what's going on.

If we go to the next slide, we're on slide 17 now. [Arm races], this is a big deal for us. We have quite a bit of [inaudible] on the Internet for multiple reasons – surveillance, censorship. You can imagine what any countries or present regime can do when they [hack] your Internet.

If we go to the next slide, on slide 18. I talked a bit just before how to block Tor networks. There's multiple ways. First of all, of course is blocking the [inaudible]. There is nine of them. They're known. And if you block them all, well nobody can bootstrap because they cannot get that consensus, that document that explains and details the [view] of the network.

It's been done in countries, in multiple countries, but there's a way to pass around that. This is why [inaudible] cache exists, that every relay can cache the consensus. But of course you need to bootstrap at some point, but you can get a consensus from outbound, from someone else, and then start to bootstrap and stuff. So this is not like an ideal way. There's also blocking all IP of the relays, as I said.

Filtering [bays] on Tor's network fingerprint. So Tor, when it connects from relay to relay, uses TLS. And inside that TLS session there are what we call the onion skin or the onion layers. So there's quite a bit of encryption, but this TLS, we try to look as much as we can as it's [HTTPS], but it's not that easy all the time. So any other thing is preventing users from finding the Tor software.

I have a few slides to show you, which is a screenshot of pages at torprojcect.org are being hijacked in different countries to show you how they prevent users [from doing that].

If we go to the next slide – slide 19 – this is an example of directly connecting users from Egypt. If you recall back in Egypt, you can see this line going down. It is when, for three days if I remember correctly, they completely shut down the Internet. You can see that drop was quite

significant and those dots show censorship [events]. Well, it's based on estimation. You can get some more information [inaudible] torproject.org [to explain that].

But this is interesting stuff where we can see how countries are actually trying to censor or there's actually movement on the Internet or stuff going on.

I'm going to show you this next slide, on slide 20, which is Libya. When this war raged in Libya four years ago or something like that, you can see this huge drop where they completely dropped. I think they even shut down the Internet for a few moments there. Then it continues to grow.

But the huge line that goes up to almost 300 users, it's usually when you have very significant political activities in the country. I have multiple graphs. I just picked some of them. But from Congo, from Ukraine, from – I had another one. I don't remember the country. Where you can see [inaudible]. You can see all of those censorship events or political events with [a rise in Tor] users.

If you go to the next slide – slide 21 – this is Iran. So Iran has been one of the more active countries blocking Tor. They did that in multiple ways. But somehow in January of 2011 they figured out quite a bit and it dropped. So almost zero [clients]. And if I remember correctly – I might be wrong here – they blocked Tor based on [TPI] of the Internet and they found a way to identify [TLF] connection. Then in my recollection, Tor just changed one of the [TLF] elements and we just

back and went through the [Iran] firewall. You can see those quite a bit with [inaudible] all our metrics. Pretty impressive.

Now, the last graph I have for this round is slide 22 where you see China. China is quite an adversary. They're pretty impressive. They blocked Tor in multiple ways. They blocked Tor… The four ways to block Tor, they did them all. And one of also the things they did is enumerate bridges. I think there's three or four different ways you can get a bridge either by e-mail – you go on the bridge torproject.org and then you can get a bridge from there. So there's multiple ways. They [inaudible] all of them and blocked them. Not only that, but now we know that when there's a [TLS] connection going out that looks suspicious or whatever, they actually connect back to you to try to fingerprint. This is one of the ways they actually are quite effective in blocking Tor. This graphs shows that in the middle of 2010 they actually succeeded quite a bit. This is an old graph, but I believe that up to 2015 it's roughly the same.

Next slide, slide 23. Those are essentially for the next two slides are pages that people get in different countries what is torproject.org. One way of stepping or censoring or stopping torproject.org would be just to not allow you to download the software.

The first one you see United Arab Emirates, Bahrain. They look friendly. There's a nice clown. Or not clown, but I mean a lady that's just telling you that it's bad and you should surf safely. This Tor project thing is not good, right?

If you go to the next slide – slide 24 – they have friendly characters telling you that it's a problem and so on and so forth. So people, when

they go to Tor project and they see that, they don't think it's censorship. They just think it's being for their own good. That's quite a bit of an issue here.

If we go to slide 25 now, what are we up against? I told you a bit about China. We have government firewalls.[inaudible] versus Iran versus China. We know a bunch of countries are trying to block it or [inaudible] move it away. Even American companies have dedicated software that blocks it.

One of the things [inaudible] is also [inaudible] provides an [option] for their users just to block Tor, and of course I think they're just blocking Tor by blocking the IPs or getting the TLS signature. But it's still something that is now being put in more and more commercial product to block Tor users.

If you go to the next slide, again, Iran 2015. You can see they blocked it quite a bit again. Every time a response to either political activities there or they just bought new hardware, or at some point their researchers a way to just filter it completely.

This arm race we're facing is a constant arm race right now. We see those censorship events happening in every country everywhere all the time. A bunch of countries [inaudible], but there are more [inaudible].

If you go to slide 27, as an anonymity network, Tor's safety comes from the diversity of the users. If you are in… Let's say you're in a country, a [inaudible] country, and you are the only one Tor users. Well, you're going to get noticed quite heavily. It's quite easy to see Tor traffic.

So the name of the game with Tor for anonymity is diversity and as many users you can. As long as there's, activity there is anonymity because then you look like everyone.

Point number two here is that 50,000 users in Iran means almost of them are normal citizens in a way because they're all doing normal traffic. And as many users as you can, well, better chances of anonymity because few users is actually easy to get to.

There is this story about a student in [inaudible] who actually used the Tor network to send e-mail to – I don't recall if it was blackmailing or whatever. But he was the only one on the [inaudible] network, so it was kind of easy to [get them]. So diversity and lots of users is actually very, very important for anonymity.

Slide 28 now. This is one piece of the puzzle. We've seen with the latest [hacking] team and [inaudible] leaks, which are those news, if some of you recall, where they sell spyware to companies, government, or whatever and one of those was actually actively snooping on Tor users. So we assume that this is a piece of the puzzle, but it's not a full bullet proof solution because if you have hardware or software that is attackable – spyware. Do you really have a genuine copy of Tor?

On the technical part, [inaudible] recently – I mean recently… I think it's been a year, maybe less. Tor is… We have created this thing which is called the [productive build]. So Tor is being built on different machines by different individuals and they all match the same [inaudible] at the end. That means it's the same copy everywhere, what we call the [productive build].

We have this great person at Tor who works on that and created the [productive-build.org] or something like that, and now [inaudible] has almost 80-85% if not 100% now of their packages being [productive builds].

This is a huge effort of Tor, so when we provide the Tor binary to people around the world, we know that is the exact copy that we think it should be.

If we go to slide 29, this is a slide from the NSA secret document that was led by Snowden. One of the things with Tor is that they don't like it at all. That was 2012, so three years ago. So that's [inaudible]. But yeah, it stinks and they don't like it very much.

If we go to slide 30, that logo was I think actually in a slide and with that quote still looking at a [inaudible]. They said that [inaudible]. We don't think that. There's multiple other network [rounds] that use different things like [inaudible] or some of the alternatives. They different business models. They also have different ways of doing Internet anonymity. But all in all, Tor was the one that was studied quite a bit by NSA. We know about [this] because of those slides. We know it has some traction.

I don't know if you guys knew about that, but two or three days ago there's this story that the FBI actually funded a university – a [inaudible] university – for [inaudible] Tor users which we fixed in July, last July. Not last July, but July a year ago. That research for the anonymity of Tor users they actually succeeded and they provided the FBI with IPs so they can do convictions.

So we know for a fact that Tor is being used by bad people. That we know. But we think it's not the main use case. I have a graph to show you about that.

Slide 31 goes to the perception. What is Tor? We fight in the news quite a bit about that on the [inaudible] Tor. Usually when someone heard a bit about Tor heard about the dark web or the deep web. We try to not use that sentence there because, for us, Tor is just Internet in a way that is differently accessed, but in a matter that protects your privacy.

So it's a huge problem perception towards the public and also towards the government or agencies that actually see Tor network as bad or not good.

If we go to slide 32, there's multiple ways to [inaudible] Tor in legal ways. ISP [inaudible], for instance. Running an exit relay is very difficult for ISPs because they get a lot of [inaudible] notice or complaints because the traffic is not all the time pretty [clear]. So that makes the ISP hate us more and more and more. That's unfortunate. We're still struggling but working it out most of the time. We have almost 7,000 relays now.

[inaudible]. That's easy. Recently – and I think still, it's still an ongoing process for people in the US. Can't go to healthcare.gov through Tor. That's kind of a shame because protecting yourself on the Internet about your health status or your health conditions or just trying to get insurance is something we should provide privacy [inaudible] and we have none right now.

Next slide, slide 33. You guys all remember that [Prominola] person, which is Snowden. Because of this [inaudible] here on top right there, and he's also advocating quite a bit for Tor. We got quite a bit of an influx of users at some point and also interests in multiple ways because this story made a trip around the world.

If we go to the next slide, on slide 34, this is a huge job in [inaudible] connecting users. The thing is that this is also [quite] linked to a botnet that was active on Tor and we quadrupled our users. We had more than that – 5 million users. So that's with a botnet, but also there was a huge Snowden [effect] after a while where people said, "Oh, what is this Tor thing? Why is this person using that?" Some interest.

So just to tell you that this line went back when we killed the botnet, but it's still at 2 million users, not 1 million, so the Snowden effect was very strong.

If we go to slide 35, I think that slide I just put it there because I don't remember. I think it was supposed to be in before. But yeah [inaudible] connecting users from Turkey, and just before September, it was the – I don't remember. I hate a note, but I think this slide is just [inaudible], so let's go to the next one, slide 36.

I'm arriving to the end of my talk. I'm just going to talk to you a bit about hidden services. So hidden services is something that you might have all heard about which is actually in the news directly connected to the dark web, which is those addresses [with a .onion].

So this is not a [classic] DNS lookup here. A .onion is an address that you can reach [a service] on Tor. The good thing with [inaudible] service is

that the clients enter the networks and will never leave it. You're reaching a service that is inside Tor network so you are not going through an exit as your data goes in [the clear] on the open Internet.

I'm going to stay on this side 36 now. We have [inaudible] and Facebook – Just recently Facebook created its own onion address. This is great because it makes a distinction between the open Internet and the dark web, or what we like to call onion space. Going to [inaudible] almost the same because you get Facebook through a .onion or a .com. For us, going to a .onion, is just a way to reach a service in a secure way, but shouldn't be seen as different from the open Internet. Yes, it has different mechanisms, but we all have different mechanisms for a bunch of stuff around the Internet anyway. So [inaudible] Facebook. We're trying to advocate for other services to provide onion addresses because it allows users to reach Facebook or [inaudible] using anonymity but also a very secure way. I'm going to explain to you why those services are much more interesting.

If we go to the next slide, slide 37. There's an address on the bottom because you have six or seven images for hidden services. As a matter to explaining to you at once, I just put one graph, but it's showing the whole thing. You have this address at the bottom and you can read about it and everything. But I'm just going to do this rough introduction of what is [inaudible] service or onion service.

You have Alice wanting to connect to Bob. Let's say Bob has a web service. Let's say Facebook. It wants to provide service through the onion space. Now, Bob's service is going to connect to three introduction points – the IP1, IP2, and IP3 you see in this picture. Those

introduction points are just a circuit through Tor. So three [inaudible] circuit as we saw prior in the presentation. And Bob's going to keep this circuit open to the IPs.

Now, Alice is going to – [inaudible] going to connect to Bob with this onion address, will be able to ask with this onion address. It's a [ash] of basically the public key of the [inaudible] and some other information. The reason is that this long address gives you a place to go in [DHT] which is what we call a hidden service directory, which are basically relays that fits the profile that is up-time more than 96 hours, fast and bandwidth, some requirements like that.

This greats a [DHT], or a ring and with the finger print of the [inaudible], of the relay that has the hidden service directory. So Alice is going to connect to [inaudible] because with the onion address she knows which [inaudible] to connect. She's going to download what we call an [agent] service descriptor which is a document explaining where to contact with Bob.

In this document, there are the IPs, the introduction point 1, 2, and 3 addresses. So Alice would connect to the introduction point, one of [inaudible] introduction point. And at that point there is a circuit from Bob that is [inaudible] and Alice is going to [ding] them. So Alice is going to tell Bob, "Please, join me to this rendezvous point."

So Alice connects to IP1, then can talk to Bob through the circuit. Then Bob connects back to the rendezvous point. And Alice connects to the rendezvous point and they both connect and now we have this whole connection between Alice and Bob. That is all through the network. That

never exists the network. They have some nice properties for secrecy and no content is actually leaked from inbound or outbound at the Tor network.

So this is roughly hidden services. Please feel free to ask questions because most of my work right now at Tor is in Hidden services. If we go to slide 38, I'm almost done here. This is as of September. Actually, no wait – way before that. But this is just September to November. We have hidden service statistics. We have two different statistics that are collected in [private] anonymous ways that tells us the amount of traffic on the Tor network.

Right now almost 900 megabits a second of the traffic on the whole Tor network is hidden services. That is around… I think our estimation is around 5-6% of the whole traffic of the Tor network. So this is why… This is one of the good reasons we use this graph when we did this work so we can show that actually Tor is not all bad traffic that goes through hidden services for marketplace for weapons or completely bad stuff in terms of human trafficking and so on and so forth. It's actually not true. We don't know what's the percentage of that traffic it is, but still it's a very small percentage of the whole Tor network.

If we go to slide 39, this is the amount of unique onion addresses we see at all times in the Tor network. Around 30,000 [at all times] that are alive. Keep in mind that onion addresses are used by services like an HTTP server, mail server, but also there's a bunch of applications. One of them is called [Rickoshare]. We can go onto [rickoshare.im] which is a chat application that uses hidden services and two people – two hidden

services – and you can just chat over Tor. That creates much more onion addresses. But at [all time], 30,000. That's our estimate.

Then if we go to slide 40, this is a screenshot of the Tor browser. This is how most of the people actually use Tor. So if you go to torproject.org, you can download this Tor browser bundle. It is basically Firefox enhanced with Tor capabilities and also various privacy [inaudible].

We have a huge team at Tor that works on that, a big team that works on that. They're amazing people and they do amazing work. This is a wonderful piece of work that is [productive build].

Finally, slide 41. This is [inaudible] on mobile, Android. I'm not sure if they finally have iOS right now, but [inaudible] to check that. But at least [inaudible] very nice. You can use Tor on your mobile phone. This is getting much more traction lately. We had some folks from Mozilla on the Firefox OS team that are working with us on getting it on mobile. Tor also has a private mode of Firefox that would be quite amazing in terms of users and reach. But basically that's how it works.

I'm going to stop there, if you notice slide 42. Of course thank you very much. I'm at 41 minutes, so that means you have 20 minutes, guys, or more than that. I don't mind to ask as many questions as you want as long as I can answer them. Thank you.

DEV ANAND TEELUCKSINGH: Thank you, David. Thanks for that very informative presentation. There's quite a few questions that attendees will be asking. First I have Glenn. So Glenn, please, you have the floor.

GLENN MCKNIGHT: Thank you. Again, David, thank you so much for your presentation. I'd like to read a couple of questions from the chat. I'd like to go back. We had a question from Alfredo Calderon from Puerto Rico. He's from ISOC Puerto Rico. His question was with regards – I'm just backing up.

It was regarding [dot.dot.go] earlier on. He asks, "Is it is true that it is advertised as good for schools and private groups?"

DAVID GOULET: So [dot.dot.go] as a service, we advertise as a solution for Google towards privacy enhancement. Because they don't keep logs, also the search requests. So yes, it is. If I understand the question correctly, yes, it is an improvement from Google. This is one of the reasons we actually add an onion address quite early, so people can reach them in a [manner] of – through anonymity network and privacy. So I hope I'm answering your question.

GLENN MCKNIGHT: The second question from the chat. Everyone, David is not on Adobe Connect, so I'm taking the liberty to read the questions that are being posted. The second question is from Satish from India. I believe he's from [inaudible]. He asked to you, David. He's been an avid fan and user of Tor for several years. I believe it is extremely valuable, that it saves lives. You mentioned last week that SMU/SEI action jointly with law enforcement. This seems to make Tor appear vulnerable. Is this a slippery slope?

DAVID GOULET: Very interesting question. Right now, all work at Tor is open source. It's a best effort. We have developers, researchers. I've been doing that for many, many years. We think we do our best work. We think our users are the most important thing. The safety of our users is very, very important. So this is why all of our development processes, research processes, are always, always, always based on user security and user safety – and also not breaking any anonymity.

Now, we will always face threats from researchers around the world, either in US government agencies or universities that actually try to break Tor because that sells a lot of papers. But also improves security of Tor.

So I would say we should continue to use Tor. That is for sure. Our work, the more we get funding, the more we get stability in our organization. The better we can address those issues, and also we can improve the services right now.

Just to give you an example, I know we have [next generation] hidden services, that with this [next generation] hidden services proposal that is being created three years ago, that attack that FBI and CMU did couldn't be possible because we are aware of issues. We still want to fix them, but it takes quite a bit of engineering effort and also research.

I wouldn't say this is vulnerable because we at Tor are very, very active at looking at the latest research, responding to threats and vulnerabilities and we spin out information to the public with our blog and mailing list and new software.

I wouldn't say you're vulnerable, but keep in mind it's not a silver bullet. That is always the case for any software.

DEV ANAND TEELUCKSINGH: Thank you, David. Thanks, Glenn, and thanks David. Next speaker in the Lutz Donnehacke. Lutz, you have the floor. Go ahead.

LUTZ DONNEHACKE: Thanks. Thank you for the wonderful presentation. I do see it coming back to ICANN that a lot of ICANN services are going to require login and real names to use the service. For instance, the new learning platform requires that you have to give your full name in order to see a simple introduction video.

The question I have is do we have a good argument to say to the people at ICANN to stop de-anonymizing the users to stop finding out who is using which service. Do we have a very good argument for how to use the learning platform without knowing who is accessing the individual files? Thanks.

DAVID GOULET: Okay. That is a very difficult question. There's lots of reasons why you would either log in with real names or not. Maybe they're good, maybe they're not. But one of the things you have to keep in mind is more a service logs real names, it creates a digital trail of what the person has been doing or is doing. This is metadata that are very useful as a reality nowadays that we're [inaudible] surveillance is everywhere.

Now, if you have the data, that means you can – if you have those data, that means you can give it back to any agencies if they have a lawful warrant or anything. But also laws change in different parts of the world. It's not something that is super stable. That means if you have the information at some point, it might not be useful at that other point. It could be used against the users or not.

My argument here would be if there's no real need for adding the names other than just, I don't know, statistics or better ad delivery, we shouldn't because we lose this privacy part. And also, this data, we don't know what it's going to be used against or for after multiple years or in the next month. Again, there could be arguments for having real names. For instance, if you want to talk on the forum, for instance. I want to know that I'm talking to Glenn or not. But it depends. I will still use this argument of no data [inaudible].

DEV ANAND TEELUCKSINGH: Thanks, David. I'm just looking to see if there are any questions. Well, actually, I have one question, David. You touched on it before. A lot of it seems to have been geared towards desktop computers and so forth. But now [inaudible] more and more Internet access is being done on mobile devices. Okay. So you have Android as one supported platform, but can Tor be installed at, say, the router level so that all of your devices that are connected to the router can use Tor?

DAVID GOULET: So your question is can you have devices – can you have, let's say, router. Then all the devices connected to that router, they can go through Tor? Is that the question?

DEV ANAND TEELUCKSINGH: Yes.

DAVID GOULET: So, yes, this is totally possible. There are multiple ways to do that. Now, just to keep in mind is the application – the Tor operates on the transport layer. So basically on TCP it makes it so that you're anonymous. The problem is every application, when they're designed nowadays, they leak a tremendous amount of information about you.

That means if you have ten devices and you go through a router, and it all goes through Tor, then all those devices will [leak] something very specific to you – about you – through the exit nodes.

Now, this was not a [inaudible] model. That was [inaudible] model, but that was not as serious as it is today, because today we know that there's a global [inaudible] over the Internet. That means when you have a global [inaudible] of the network, that makes it much more difficult for anonymity network to operate in a way that we have certain percentage of guarantee that you are fully anonymous. In this case, if you have multiple devices going through Tor that all [leak] something about you at the exit nodes because you've got to exit the network at some point.

Keep in mind that it's something that could make you stand a bit more out of the rough. So this is why Tor browser, for instance, is trying to make you look like any other Tor browser user, so you're not fingerprintable.

DEV ANAND TEELUCKSINGH: Okay. Thanks, David. I see Glenn has a hand raised. Glenn?

GLENN MCKNIGHT: Hi, again. I have another question from Satish Babu. His question is, "How scalable is Tor? If many more users start using it and the number of nodes are roughly the same, will the performance start degrading? Also, is there anything [possible] [inaudible] can do to save this situation?"

DAVID GAULET: Very good question. Scalability is always something we have in mind. Right now, we are working actively on scalability of [inaudible] services. If you recall from this presentation, connecting to a hidden service is actually very, very intense work on the network and also in cryptography. You build three different circuits. Then you connect to this other person. It's a lot of cryptography and a lot of bandwidth being used just for connecting.

We are working on trying to scale this in multiple ways. We created onion balance out of [inaudible] privacy program, which is basically a [GSOC] program, [Google Summer of Code]. And onion balance [inaudible] services. Facebook is actually very interested in that and

they're using it, beta testing it. [Low balance] is one onion address through multiple coordinates.

But if we go to the network part, if the relays right now stay the same, if you look at the graph on metrics, we still have a bit of a gap between what is used and what is [inaudible] to use, but it's always a problem.

If, let's say, we go to 4 million users with those 6,000 relays, chances are that the performance of the Tor network will degrade. This is the most important part is that, as we're getting more users, we should get more relays – fast relays that allows from exit to an entry point powerfully fast.

So scaling is an important question for us. Every time we add some security component to Tor – for instance, just to give you an example, we're working here to add padding on Tor network. So traffic [correlation] or traffic confirmation attack looking at two different parts of the network and say, "Oh, you are David." It gets much more difficult. But then it adds [loads] to the network. So this is always in our research and development processes. Can the Tor network scale with what are adding? And right now I cannot give you a firm yes or no, we are scaling or not. It's a constant ongoing process. Thank you.

DEV ANAND TEELUCKSINGH: Thank you, David. I'm just seeing if there's any other quick questions. Let me just ask a quick poll for all our attendees, just to get some idea of this. How many people are using Tor? You can use your Adobe Connect to say yes or no. Go ahead, David. I suspect you are, of course, a very heavy user of Tor.

DAVID GOULET: So the question is how many people do use Tor?

DEV ANAND TEELUCKSINGH: Well, it's more for a poll for the audience. But you can answer that question while people are responding.

DAVID GOULET: Oh, for the audience. Okay.

DEV ANAND TEELUCKSINGH: I'm seeing that there are two persons that indicated that they use Tor. It's about 14 or 15 persons attending the call. Three persons I see now. Excellent. Okay, just seeing if there's any further comments or questions. Going once, twice. Okay.

David, thank you so very much for this presentation. It was very, very informative. I think perhaps what we are going to be doing is we'd be putting the slides up on the wiki and perhaps still maybe follow-up questions. Via Glenn, we could probably relay these questions to you if there are any follow-up questions. And perhaps if the Technology Taskforce has further questions, we may want to further contact you. We will keep in touch.

DAVID GOULET: Yes, please do.

DEV ANAND TEELUCKSINGH: Thank you so very much.

DAVID GOULET: Ask questions, put those public. It's my pleasure.

DEV ANAND TEELUCKSINGH: Great. Thanks so much, David. Well, we are four minutes past the top of the other. We did have other agenda items. I'm just going to ask for perhaps an additional five minutes just to quickly talk about possible immediate work items, objectives, policy [TF] between now and the next ICANN 55 meeting.

So one of the things that we are going to be looking at are conferencing solutions. Now, as you may recall, the Technology Taskforce has been [inaudible] conferencing solutions which started off in 2013 when ICANN is, at one point, considering switching from Adobe Connect as a conferencing solution to something called Lucid meeting. And the discussion right then when that was happening was that, well, if they were going to switch, we should probably have some say as to what type of conferencing solution we want.

That started a process where we started reviewing conferencing solutions. When ICANN seems to back away from switching from Adobe Connect, we also kind of stopped reviewing conferencing solutions. But I should say that at the Dublin meeting – and perhaps, Glenn, you could step in onto this because you had a conversation with ICANN staff about this.

Some of the ICANN staff members talked about, well, we feel that we are using Adobe Connect, but we are looking for alternatives, Glenn, do I have that summary correct?

GLENN MCKNIGHT: Sure. Yeah. I did talk to Paul [Hoffman] and he will be on our call either January or February and will be focused on comparison of best practices. We did a call last week on [Flick] meeting which was a very successful call with those who attended. If there's any conferencing tool the community can recommend that's going away from Adobe, then what is he is saying is we are looking at other alternatives. So please give us a solution. We'll be happy to do a separate call that showcases alternatives. We want to make sure that accessibility works. People on every platform, on Linux, [DSD], mobile. We want to make sure it works in countries that have slow bandwidth. So we have some challenges.

Please, if you have any solutions, let us know. We do have on our site a comparison of different tools and I also have a slide show that I provided which is a little bit dated now, but it's still a very valuable tool developed [with IEEE].

Back to you, Dev.

DEV ANAND TEELUCKSINGH: Thanks, Glenn. Just to answer a question from – well, I don't know who, but from C2 in the chat saying moving something to [inaudible] not Adobe would be good. We do have some sort of feature comparisons. We did come up with a set of features that we would like to see in

conferencing solutions. We probably wouldn't have time to go into every single feature, but I'll put some links on the Technology Taskforce list, what are the core features that we're looking for – multi-platform, accessibility as Glenn mentioned, those types of feature. Must have the ability to see participants, be able to raise hands so we can have a speaker queue and so forth.

We do have a list of features that we're looking for. Again, please post in the chat or on the Technology Taskforce list or some ideas that we want to look at. So that's one of the things that we want to focus on between now and ICANN 55, which is Marrakech which is in March. We're in mid-November, so that's really December, January, February.

One of the other things we probably might want to look at is, well, we have two projects that were inspired by the work in the TTF and that's the e-books and captioning. Perhaps we can devote some more attention to those projects. Those projects were approved by ICANN and they are now two formal projects that have now been taking off since coming to Dublin and now afterwards.

Glenn, you want to quickly just mention the e-book idea?

GLENN MCKNIGHT: Sure. Maureen and I have been working hard on looking at a lot of different tools from [book write] to [inaudible] and using Caliber for the conversion. We're testing it on Kindle and iPad. We will have two e-books for demonstration very soon. We'll do one of the [inaudible] in part with the captioning pilot at one of our calls in the future. It's actually been an interesting project because what we have is a history

of doing webinars with capacity building, but that's it. Only a few people attend. This will give us beyond the webinars. The idea is to provide it right across the eco-space. It's going to be an interesting little webinar.

DEV ANAND TEELUCKSINGH: Okay, great. So thanks, Glenn. That's probably the second part of what we need to look at. In terms of the ATLAS II recommendations, I know that – I think there's going to be some next steps happening going to ATLAS II implementation in terms of how we are going to be dealing with those recommendations because a lot of work has been done by the Technology Taskforce and other groups on these recommendations.

We probably will review them again and [inaudible] possibly see some interest in would be the actual – looking at other participatory mechanisms that's in one of the recommendations.

For example, one of the recommendations that somebody had pointed out was to look at things such as Liquid Feedback and so forth. We haven't really had a chance to really look at that in any form, as such. I don't know if it's… I know Jimmy has had some experience on Liquid Feedback and I think he made a recommendation that we shouldn't do it, actually. It was probably one of his recommendations. But we do need to at least understand it, so at least we can say, "Okay, these are reasons why we can't look at it."

We probably don't need to do that much more work on the ATLAS II recommendation because I think we really have it well at hand, but of course anybody in the group can disagree and suggest we need to do more work on it.

Any other proposed ideas for the TTF? Going once, going twice.

Actually, there is one more issue I want to bring up, actually. That is we also want to catalog some of the technology issues that the At-Large community has. I started building up that on our Technology Taskforce workspace to track the technology issues.

For example, for Adobe Connect, there was an issue regarding e-mails primarily from users that are using yahoo.com addresses. They were being bounced from the mailing list. So I want to track those types of issues that the community are facing and then raise it for the staff so that we could try to see if we could come up with – well, identify the problem, come up with workarounds or solutions in coordination with ICANN staff. So that's another immediate work item for the TTF between now and ICANN 55.

One final thing. I want to spend just three or four minutes. Alternative times for Technology Taskforce calls. Typically, the Technology Taskforce meets usually on a Monday, usually the third Monday of each month at 15:00 UTC. Now, what some persons – primarily for the Asia-Pacific regions – have been saying that this is probably a very inappropriate time to attend the Technology Taskforce calls.

I just want to raise it up for the group. Should we have alternative times for Technology Taskforce calls? Do we want to alternate Technology Taskforce calls as a trial just to see? When I say a different time, I'm talking around – I'm trying to remember what was the suggested time. I believe it was something around 21:00 UTC or 22:00 UTC. That will put it early in the morning in the Asia-Pacific region. Does anybody have any

immediate thoughts or comments on that, as to whether to try for alternative times?

I'm seeing somebody typing in the chat. I'm also seeing Maureen typing in the chat. Okay, let's try alternative times. Maureen is typing. While they're typing, Glenn, go ahead. You have the floor.

GLENN MCKNIGHT: Yeah. One of the things, folks, we talked about is because we have so few… And I do appreciate Maureen being on the call today. We all know it's probably 3:00 in the morning for her. We're looking at definitely reaching out to APRALO and having a call at a time that's very good for them. Myself or Dev, I'd be happy to chair the meeting. We need a co-chair for that call in the APRALO area. If we can find… We'd like to be on the next APRALO call to bring this up as an issue. But I think there's lots that the APRALO people can participate with. Yeah, we'd be happy to accommodate other time zones. Back to you, Dev.

DEV ANAND TEELUCKSINGH: All right, great. Thanks, Glenn. [inaudible] what you are saying. I think we will definitely welcome changing to an alternate time. Probably just rotating the times. We could bring more persons, especially from the Asia-Pacific region. We do have a lot of people that are now listed from the Asia-Pacific region, but they're not on the calls. Perhaps by switching the times, we could get more of those persons in.

Perhaps what we can do is probably raise it on the APRALO. Maybe we'll then have a Doodle poll to pick an appropriate time and see where we can go from there. For Maureen right now, it's 5:00 AM.

Any other comments or any other business?

Just to mention, also, regarding interpretation and so forth, for any other working group, typically how interpretation is handled – and staff can correct me on this – usually three or more persons have to request interpretation before a session, and we have to at least give 72 hours or three days' notice before, because obviously it takes time to organize whether the interpreters would be available for that date and time. So we need at least three persons requesting a translation and for a particular language channel. That's something to keep in mind.

I know that there was some comment that we should have – I think it was from Internauta Venezuela that this session should have been interpreted. But what we are going to do is translate the transcript from this call would be translated so you'll be able to read the session in Spanish and French.

I'm seeing no further comments or questions. And I know we have gone beyond 18 minutes past the hour. It was really an informative session. Thanks again to David on his presentation on the Tor projects.

I would like to thank everyone. Do interact with us on the mailing list and this call is now adjourned. Thank you and have a wonderful day/evening/morning. Bye.

TERRI AGNEW: Once again, the meeting has been adjourned. Thank you very much for joining. Please remember to disconnect all remaining lines and have a wonderful rest of your day.

[END OF TRANSCRIPT]