SAN JUAN – Tech Day Part 2
Monday, March 12, 2018 – 13:30 to 15:00 AST
ICANN61 | San Juan, Puerto Rico

UNIDENTIFIED MALE: Tech Day, Part 2. Monday, March 12th, 2018, from 1:30 to 3:00 P.M. Room 209-BC.

EBERHARD LISSE: Okay. Good afternoon. If you can all settle down, then I don't have to call you by name, for the record.

Maybe I should call Robert Martin-Legene by the name for the record for not sitting down.

Okay. Thank you very much for coming back after lunch. This morning, we had well over 100 in attendance, which is quite good. Usually we have more than the ccNSO has on their Members Meeting anyway.

First in the afternoon is Pablo Rodriguez from .pr. He will talk a little bit about the situations they set up and whether and how it was affected by the recent climatic events.

PABLO RODRIGUEZ: Thank you, all, and good afternoon. We are honored to have you all in Puerto Rico. We greatly appreciate your support. Therefore, welcome to Puerto Rico. We're very, very glad to have you all here.

This morning, I would like to quickly go through some of the accomplishments that we have had throughout the years. One that we are especially proud of is the fact that .pr was the second ccTLD right after Sweden in adopting DNSSEC. That is a fact that not too many people know. We would like to stress it and make it a point that indeed we did. We presented this in 2006 in Sao Paulo. We have actually some of the people that helped us make that implementation at the time.

Some of that implementation helped us by promoting the adoption of DNSSEC. It helped us by also inviting and inspiring other people to adopt other technologies in addition to that.

I am going to skip some of that information to go onto what really brings us here. Our experience has helped us learn some very important lessons. The idea is that our region is extremely vulnerable to natural disasters. Our region doesn't get one natural disaster. Sometimes we get two or three rather close to each other.

Let's take a quick look at that – whoops. Can we run them? Thank you.

UNIDENTIFIED MALE: President Trump declared an emergency today in Puerto Rico. Residents were warned to brace themselves.

UNIDENTIFIED FEMALE: [Speaking in foreign language]

UNIDENTIFIED FEMALE: Maria is expected to become a hurricane later on. Its path is actually following the path of Irma.

UNIDENTIFIED MALE: Puerto Rico [Speaking in foreign language]

UNIDENTIFIED FEMALE: [Speaking in foreign language]

UNIDENTIFIED FEMALE: There are no lights. It is dark. Entirely dark. And apparently it could take months to get the power back on. Just think about that for a second in the modern world: months with no power.

PABLO RODRIGUEZ: Months with no power. It wasn't only no power. When you don't have power, your water pumps don't work, so you have no

water, either. You don't have any telecommunications. Besides not having any telecommunications, you don't have your elevator. The most mundane task – taking the garbage out. I live on the 18th story of a 21-story building. Taking the garbage out is a major task. It's walking down 18 stories, walking up 18 stories. Going to get groceries, bringing a couple bags: it's 18 stories up, 18 stories down. There is no gasoline. There is no radio.

I remember turning on the radio and listening to a baseball game I said, "Who the hell is playing baseball right after a hurricane?" Guess what? Our spectrum was so empty that we were capturing a baseball game from Venezuela. There was nothing. Silence.

That's when you begin to think about this. Are you ready for this? Are you prepared for this? I must tell you in advance that NIC.pr was operational prior, during, and after the hurricane. We were 100% operational.

But the idea is that, less than 24 hours before the hurricane, Mexico had an earthquake. Hundreds of people died in Mexico City. A significant number of buildings collapsed. Right after that, a major Category 5 hurricane passed by Puerto Rico.

Guess what? Prior to Hurricane Maria, we had Hurricane Irma on a Wednesday. That was a Cat 5 that destroyed a big portion of our infrastructure. We were already limping.

Eight days later, Hurricane Maria passed by and dumped so much water on us, on already-saturated grounds, that those mountains simply melted away, taking much of the telecommunications with them.

Let me show you something. Kim, is there any way you can help me? Show me the – that's the one. Thank you.

Now, let us take a look at that picture. The upper picture is what Puerto Rico looked like prior to Maria. It was taken from the International Space Station.

The second picture was taken about seven or eight days ago by astronaut Joe M. Acaba, who happens to be of Puerto Rican descent. They just landed in Kazakhstan not too long ago. When he took pictures of Puerto Rico from the International Space Station, that's what it looked like.

As you can see, the most lit area in that map is where we are. That's a metropolitan area. That's the San Juan area. Look at the rest of the island. It's barely lit. Therefore, you can imagine how little restoration of the service is available. The electrical service is limping terribly.

In addition to that, the telecommunication services is also limping. To talk about the Internet, to talk about anything else,

is redundant because we need to get this going first before we can start moving onto bigger and better services.

Let me show you. That's what some of our telecommunication towers look like. We've had nearly 1,630 or 1,640 towers, of which, after the atmospheric phenomena, we end up with about 300-200 of them. The rest were completely destroyed, demolished – literally ripped from their cement bases.

Look at this. These are 40-foot pole in concrete. Their core is steel. That core, when you look at it, was completely twisted.

These are telecommunication towers whose wires weigh tons. To raise them, we need helicopters. We had Chinook helicopters helping us to raise this type of towers. It was barely possible.

Over and over and over again, we experienced that type of catastrophe. We talked about disaster recovery planning, but most people tend to think about one disaster. How about two? How about three? How about getting hit by an earthquake, a tsunami, and a hurricane? Is that possible? It appears so. It appears that that can be the case.

So we used the domain name Status.pr, and we use that as a tool to provide complete transparency and information to the people in Puerto Rico and throughout the world. It was published throughout CNN, and it was the official website for

everyone to come and get information on what was happening in Puerto Rico. It's still available and still operational.

At the moment that I made this presentation, our telecommunication service is back on at about 98.5%. We continue to make improvements, but as you saw earlier in a previous slide, the restoration of electrical services is still very slow. We still need to do a lot more and a lot better.

This is one of the problems that we have. Back in the '70s, Puerto Rico bet on becoming an oil refinery hub. We developed a lot of the electrical production of electricity. It was a place in the southern portion of the island.

However, the northern portion of the island has the highest use and consumption of electricity. So after the hurricane, we ended up with production in the south, consumption in the north, and nothing in between. Trying to get the wires across our topography, although we are a very small island – only 100 miles long by 35 miles wide – was a major task, and it continues to be a major challenge. So that is part of that slow process in making the restoration of that electricity.

Now, what did we do? Well, this isn't the first time. We are not alien to disasters. Many of you will remember the disaster that occurred in Japan in 2011, where they experienced a 9.1 earthquake. Right after that, they had a tsunami. Mighty Japan

was [uncommunicative]. Their grid was destroyed – a significant percentage of it – and many, many people were unable to communicate.

At that time, we knew that we should do something to help out, to help protect the cyber real estate of domain name users in Japan. We knew we had customers, registrars, in Japan. So immediately we set out to search our databases and see how many domain names we had where the country of origin was Japan or in the vicinity of those who were affected by the earthquake and subsequent tsunami.

Well, fast forward that. Seven years later, those lessons learned from the Japanese event, where we set out to identify those domain names and then immediately extended their expiration date to avoid those end users losing their domain names, we applied on us. We began searching our databases for domain names whose country of origin was Puerto Rico or the Virgin Islands or Dominica or Antigua or Bermuda because these were all the areas that were affected by either Irma or Hurricane Maria.

We identified 626 domain names, and immediately we extended the expiration date. To this day, for some of those users who have not responded, we are still waiting for them. We are not deleting those domain names.

What's the point here? The point is that, when a natural disaster occurs, for those of us who are registries and/or registrars have a responsibility into searching our databases and trying to see if any of those domain names that exist in our databases are coming from those countries and/or territories that were exposed to that natural disaster or were affected in any way. Then we must do something to protect them.

What did we learn and what do we continue to share with our people? That the Latin American and Caribbean region is prone to natural disasters. We have seen it. As I mentioned earlier, Mexico experienced an earthquake less than 24 hours before Puerto Rico was ransacked by a hurricane. Irma had affected terribly the southern portion of the Caribbean region. We continued to help out those people, and a week later we get hit by our own hurricane.

Am I responsible for my brother? We believe so. We believe we are. I believe we are. So we did. We took all the measures that we could to protect the people that could have possibly have been affected by these hurricanes. I continue to push so that every single one of us can develop policy or help out to develop policy that protect the cyber real estate of those users in each of the registries and registrars throughout the world.

This is not something that occurs only in the Latin American and Caribbean region or in the Pacific Rim. We have seen very recently the flood in Paris. We have seen the climate change in so many different areas throughout the world. Chances are that, at some point, each one of our regions will be affected by some type of natural disaster.

Will we be prepared for that? Will we be prepared to manage more than one catastrophe, more than one disaster? That's what we need to do. Disaster recovery planning is extremely important. In addition to that, we need to have the sensibility to know that, even if a disaster occurs on the other side of the planet from your perspective, chances are you may have customers there. What are you going to do to protect them? Do you feel that you are responsible or not?

Well, I feel that we are, and therefore we [inaudible]. That said, I'd like to thank you all for taking the time to listen to me. I hope that some of these experiences can be of help to you and give you some food for thought and that some of that can be put into action and inspire you, if you have not already began thinking of this disaster recovery planning, to engage in that activity because it will make the difference between having a registry or registrar, having a company, or not having one.

Thank you very much.

EBERHARD LISSE: Thank you very much. Excellent presentation. I'll allow one question just now. What I want to say from the rostrum again, abusing the prerogative of the Chair, is that we shouldn't just look at disaster recovery or disaster mitigation. We should like it from an aspect of business continuity. You lost Professor Moreno a few years ago. My colleague just died over the weekend. Fortunately, we are elderly people, so we had some idea that that could happen. We had already put structures in place. But this is probably something that especially younger people at smaller companies don't really look at.

For example, I've got eight staff members. It's probably not a very good idea to let all three technical people go on the same plane when there's bad weather or something. Little things like this.

For that reason, I'm very keen on having, in Barcelona, a session from the technical but also from the organization side on business continuity.

Patricio is next. He had a similar situation during the earthquake. I remember very fondly the remote presentation he gave in Nairobi when he was not even there and he could do it via Adobe Connect.

PATRICIO POBLETE: Thanks. Patricio Poblete from NIC.chile. It's true. We're not strangers to disasters ourselves. I find it remarkable, what you said: that you were able to maintain operational continuity throughout all these disasters.

But in addition to hearing about what you did, I'd like to know more about how you did it. For instance, with no power, no telecommunications, and so on, were you able to operate from here, or did you have sufficient resources abroad that you could use?

Also, operational continuity: of course, I understand that refers to domain name resolution, but were you also able to take new registrations to the invoicing and so on?

PABLO RODRIGUEZ: Thank you, Patricio. It was a pleasure. Well, we had two co-locations. We had one called Critical Hub, and we have a bunker at AT&T. Both of them had generators, and our systems were running from both locations that were redundant and were mirrors to each other.

By the fifth day, one of them – the Critical Hub one – went down, for lack of –

UNIDENTIFIED MALE: [inaudible]

PABLO RODRIGUEZ: Yes. For lack of diesel fuel. Imagine, we had no diesel, no gasoline, no nothing running available. The AT&T hub had diesel for five days. We continually ran from that, and then we were praying and hoping that somehow we could get our hands on diesel. But our governor made sure that, once he got his hands on diesel fuel, hospitals and telecommunication companies were priorities. That's how AT&T got their diesel and we were able to continue. It was completely transparent to everyone throughout the island and throughout the world.

So we continue to support domain name registrations, answer e-mail, and renew domain names. It was completely transparent to our users.

Thank you.

EBERHARD LISSE: Okay. Thank you very much. Impressive presentation. Keep it up.

PABLO RODRIGUEZ: Thank you. Thank you very much.

EBERHARD LISSE: Warren Kumari from Google or in his IETF hat or whatever hat he's wearing today, if any, is going to speak about KSK Sentinel.

WARREN KUMARI: Let me just check if the forward and back works – yes.

When I started making this presentation, I thought it was about a fairly large problem. But after just seeing this last presentation, it kind of resets your expectations of what's a major issue and what's not.

Anyway, this is going to be a quick presentation on the KSK Sentinel. It's myself and Joao and Geoff. There's Geoff.

What's the problem we're trying to solve here? I'm assuming that most people here are fairly familiar with DNSSEC, or at least somewhat familiar with DNSSEC.

Okay. For those who are not, some of this will be clear. Some of it won't. So what's the problem we're trying to solve? Well, we'd really like to roll the DNSSEC trust anchor. That's the root KSK. Users who are using a validating resolver that do not yet know about the new KSK will simply break. All of their DNS resolutions will just stop working. They won't be able to resolve any names.

They won't be able to reach the Internet. They will basically be dead in the water.

More worryingly, we don't have any real way of measuring who currently has the new KSK, who doesn't have the new KSK, or who's just relying on resolvers with new ones or old ones. This just isn't something we can measure, so we don't really know how many people will be affected when the KSK roll happens.

For people who have been coming to these for a while, they might remember that there was this thing called RFC8145. It is supposed to provide reporting on this. I was one of the authors of it. It seemed like it would be great. It does actually provide some reporting, but what it reports on is actually what the resolvers will see. So this provides reporting from the resolvers to the root servers which particular key they have.

This seems like what we wanted to measure, but once it was deployed, we realized there were more things that we wanted to know. For example, I personally have a validating resolver in my basement for various long and convoluted reasons. It does not have the new KSK. However, it doesn't really matter because nobody's using it. There's a resolver sitting in a basement. Nobody queries it. Does it really matter if it doesn't have the new KSK? Nope, because when it breaks, it's not going to affect anyone.

The RFC8145 stuff does, however, make pretty graphs. It created this graph, which shows that not all of the resolvers have actually been upgraded and not all of the resolvers have the new key. When this graph was made, there were around 5% of validating resolvers which did not have the new key. For some reason, that number has increased to around 30% of validating resolvers that do not have the new key. So things have gotten worse.

But, fundamentally, this make pretty graphs. It doesn't actually tell us the most important thing, which is how many users will be affected. It tells us that some set of resolvers will die, but it doesn't provide any information on how many users or which particular ones will be affected. So we don't really know what the impact of the key roll will be.

We have this new thing which we think will actually answer these. It requires a small resolver update, and more importantly, it allows anybody who can publish something the DNS to be able to actually set up some measurement service and understand how the key roll is going.

More importantly than that, it allows user to actually test this themselves. So this will allow users to have a look and understand what the results of the key roll will be for them.

The way it does this is that it's a simple change to the DNS protocol or DNSSEC protocol. The very last thing that a resolver that implements this will do, just before sending the response back, is have a look and see if the leftmost label in the name contains kskroll-sentinel-is-ta-[key].something. If it does have that, and if the resolver has that particular key ID, it sends back the answer normally. If it does not have that key ID, it converts the answer into a SERVFAIL response. SERVFAIL is the generic DNS response for "something went wrong when I was doing things and I had an error."

The other rule: if the leftmost label starts with kskroll-sentinel-not-ta-[key].something, then, if the resolver does not have this particular key, then it replies normally. Otherwise, it answers with a SERVFAIL.

So these are the only two changes that need to happen, and a number of resolvers have started actually deploying this.

The reason we have both the is-ta and not-ta is that we can determine which particular resolvers actually support this mechanism. There's an example of that in a second. These slides work better when they're presented on the web and each line comes up one at a time, but I'll go through the scenario.

So, I'm a validating resolver. I also support Sentinel, and I have the new KSK. The new key ID is 20326. I'm sitting there and I'm

running along happily, and I get a query for invalid.example.com. This name happens to be DNSSEC bogus. It's signed with an incorrect key, so I do what I would normally do: I reply with a SERVFAIL.

A little bit later, I get a query for kskroll-sentinel-is-ta-20326.example.com. I do all of my normal DNSSEC work, all of my normal DNS validation. I resolve it. It happens in this case to resolve to 192.0.2.23. That's just an example address. It resolves to some address.

Just before I send back the response, I see that this particular query started with the magic string, kskroll-sentinel-is-ta. I check and I see that the key ID that it's asking about, 20326, is a key ID that I have. So I do my normal response. I just answer, "Yeah. Here's your answer. It's 192.0.2.23."

A few seconds later, I get a query for kskroll-sentinel-not-ta-20236.example.com. I have the key ID 20326, so I use the second rule in the serve set of changes. I convert the answer into a SERVFAIL and I send that back. So it's because I don't not have 20326 that I send that back. For The Simpsons fans, that's like saying, "I'm not not licking toads."

So now we've added a bunch of complexity to DNS. Why is this helpful? The plan is, once this is implemented, we ask users to go off to a webpage. On the webpage, we have three different

images. One of them is invalid.example.com/fish. One of them is kskroll-sentinel-is-ta-20326.example.com/kitten, and one of them is not-ta-20326.example.com/puppy.

We ask users, "Do you see a picture of a fish?" If you see a picture of a fish, then that means you are able to somehow resolve the invalid.example.com/fish name.

The only way that you could have resolved that name is if you're not doing DNSSEC validation. It's an invalidly-signed name. If you are somehow able to resolve that, that means you're not doing DNSSEC validation.

That also means that the key roll doesn't affect you in any way. You don't need to worry about this. You can just ignore the problem. There is no issue for you.

If you see a picture of a kitten and a puppy, then that means that is simultaneously true that you do have and you don't have that trust anchor. The only way that that could be true is if you don't yet understand this mechanism. You're able to resolve both of these names. You can't both have that key and not have that key. That just means you don't yet know how to do this. Unfortunately, that means that we can't actually tell you if you'll survive the key roll or not.

If you're able to see a kitten and only a picture of a kitten, then that means you've got the new key. It means that you have 20326 and you don't not have 20326. So that's how you're able to resolve the kitten and not the puppy. You will be just fine.

If the only thing that you see is a picture of a puppy, then that means that you do not have 20326. You were not able to resolve the kitten. You were able to resolve the picture of the puppy. That means that you are in danger when the KSK roll finishes. You will not have the new key, and all of your DNS will break.

Are we seriously going to do this with pictures of kittens and puppies and asking people to please have a look and tell us what happens? I wish that we would. That would be really awesome. Sadly, no. We will instead send people a big blob of Java script.

What this Java script does – I'm not sure if you can read it on the screen – is, basically, that test before, on the previous grid. It makes your browser load three different resources. One of them is invalidly signed. One of them is is-ta-20326. The other one is not-ta-20326. I just realized I have a typo on my slide. Then there's some Java script, and it has a look to see which one of the resources it was able to reach, and which ones it was able to actually render. From that, it will tell you if you can or cannot survive the key roll.

If you actually try that from this particular network, you're behind DNSSEC validating resolvers. They do not yet support this update. Eberhard is busy doing it as we speak. What you will get is a page that looks like this. It will tell you, "You're using a legacy resolver. We cannot tell you whether or not you will survive the key roll." But once these resolvers are upgraded, you will hopefully see, "Yeah, you will be just fine."

I think that this is my last slide, other than a discussion I had with pictures of kittens and puppies and stuff. So you didn't get a picture of kittens and puppies. Here's the gratuitous kittens. I don't know why they're gardening, but they are. So here's your kittens.

Questions? I think many people might have seen this at the DNS-OARC presentation, so we might not have as many questions [as expected]. Let's see.

EBERHARD LISSE: Any questions? Don Hollander has one.

DON HOLLANDER: Don Hollander from Book Haven in New Zealand. In the past, Geoff has run these very clever advertising-based sensors to see – have you done that with this?

GEOFF HUSTON: This is Geoff Huston from APNIC. I was a co-author of this draft with Warren. It's been a real problem with this roll of the KSK. The way the DNS currently works and the way resolvers currently work is that it's impossible to test the preparedness of the new KSK prior to it coming in, unless we change something. Something has to change. We either have to change the DNSSEC protocol itself, which in turn means altering the behavior of every single validating resolver, probably by bringing up a golden name in the root zone that's only signed with the new key and etc. It gets horrendously complicated very, very quickly.

Or – this is slightly simpler. It doesn't change the protocol, but it must change the action of validating resolvers, that they take what would normally be a validated secure response and, under certain circumstances, depending on the query name, and alter it from, "Yeah, that's okay," to SERVFAIL. So it's basically a downgrade from, "I'm not going to tell you the answer under certain circumstances."

We can do this with a mass [inaudible] mass testing. That's fine. You can do it yourself, as Warren showed with his page. So it's not exclusively a test that only I or someone else with golden keys can handle. Literally anyone can set this up and test wherever they want to test. So that's the good news.

The bad news is that resolvers have to change. There's a certainly a conversation with the folk who build resolver code to say, "You need to put this in your code. If you find this particular sentinel's label, change your behavior."

Folks who run resolvers – ISPs, etc. – should be keeping up with releases. If you're running a big resolver and the version is five years old, you're in deep kinds of stuff, no matter which way you look at it. You should be keeping up to date, but it does take time for this to come out.

Will we have answers? If everything goes according to track, by October of this year, which is the current plan day? I really don't know. It's up to a whole bunch of variables that I have no control. But literally, in thinking around this problem, there's no other way to do a reliable way of mass measurement that we can see that doesn't entail pain of this order of magnitude. There's just nothing we can do any other way.

Thanks.

WARREN KUMARI: Actually, following up from that, the KNOT resolver has already started implementing this. Some other resolvers are planning on implementing it, or we're talking to them about implementing it.

Once they have implemented it, there'll be some delay until they get deployed, and then, yeah, Geoff's planning on some [inaudible].

GEOFF HUSTON: And it would be great if someone like ICANN can encourage these resolver vendors to actually do this quickly. Then we'll just go and kneecap folk who run resolvers to actually implement the latest version of the code.

EBERHARD LISSE: Any other questions? Anything from the remote?

Okay. Over to Martin-Legene.

ROBERT MARTIN-LEGENE: This is Robert Martin-Legene from PCH. This is an RFC draft for experimental – [what's the] plan? Because the problem we're having now with the root we're going to probably keep seeing for every five years.

WARREN KUMARI: Yeah. It's currently a DNS-adopted IETF draft. We're hoping that it's going to go to working group last-call real soon now, possibly in London. Shortly after that, it'll be published as an RFC.

However, we think it is largely stable enough now that people can already start implementing it. We're not really expecting large changes, so resolve operators should be able to start implementing. And many have, based upon the current spec.

So, yeah, it would be nice if it went a little faster. It's probably on track to be one of the fastest DNS op-published RFCs in a long time. Not that that's saying a huge amount. So, yeah, that's the status currently.

UNIDENTIFIED MALE: [inaudible]

WARREN KUMARI: I can't remember if we put informational or standards track? Can you remember what we…

UNIDENTIFIED MALE: I can't remember either. [inaudible]

WARREN KUMARI: Yeah. We probably put proposed standard on, but I'll have a look while Viktor asks.

VIKTOR DUKHOVNI: Right. My name is Viktor Dukhovni. I also run a survey. I'm well-aware that my survey is heavily biased to measure just the things that I know how to measure. I can't measure the things that I don't know how to measure, so it doesn't report on them.

This one, too, seems to me that'll only measure the resolvers that upgraded to implement it. Those might very well be the ones that might as well damn well deploy the new key. Why wouldn't they? In fact, they should deploy the new key and maybe skip this step. Well, perhaps not. For the future.

So the question is, how will you get visibility to the folk who really need the help who are still stuck in the 2010 key?

WARREN KUMARI: What this does that's different is that it breaks the world up into four sets of people: those who don't validate, which we can detect, those will be fine because they have the new key, those who will die because they do not have the new key, and those who are just legacy and aren't upgraded. So we can classify people into one of those four buckets.

The people who are legacy and we can't tell any more is our error margin. So it provides data. Yes, it would be great if anybody who upgrades and then stores the new software that supports also installed the new key. However, sadly, for many

implementations, upgrading keeps the same config file. With BIND, at least, there's a stanza in there. You can either have managed keys, which say, "Do all of the key roll magic automatically," or it has trusted keys. If it says trusted keys, it says, "Use this key and only this key."

So, yeah, potentially a reasonable thing for resolver implementers to do would be to just ignore trusted keys and pretend that it said managed keys or something. But what this will do is let us actually know the error bars. It will deal with the fact that some resolver software doesn't yet understand that 5011 is probably – automatically rolling the key is a good thing.

The other issue with that is, I think what we might have seen – why we've gone to 30%; people who only have the old key. The process for rolling key says, "When the old key introduces the new key, you have to wait 30 days before you start trusting it." Many people have machines that are virtual machines or run [Docker] or something similar and they learn the new key, but they're never alive for the full 30 days, so they never actually finish the key roll.

So, yeah, it would be great it updating things fix it, but sadly not. Geoff has a comment, I'm guessing.

GEOFF HUSTON: Yeah, this is an additional comment. Two things. Firstly, Viktor, we told ourselves exactly the same thing over a year ago; that folk doing DNSSEC validation turned it on deliberately, that obviously they were going to follow a key roll. This whole thing about looking for folk who were filing was kind of irrelevant because there really wasn't going to be an awful lot of folk there.

Part of the reason why it got postponed was that that signal was completely unexpected. We couldn't understand what was going on. So there is some problem going on there with RFC5011 following a key roll.

However, understand another thing, too, which is subtle but important. 8145 is about resolvers. This sentinel is about you as a user and your cloud of DNS resolution. The sentinel will not identify recalcitrant resolvers. It doesn't try to peek into the cloud of the way DNS is resolved. It simply says, "As a user, with those lists of resolvers that you have in [etceterasolve.com], are you going to get affected by a key roll?" Are all of your resolvers validating? Because if one of them doesn't, you're okay because, when you get stranded, the SERVFAIL will send you to the one that doesn't validate. If all of your resolvers validate, the test goes on to say, "If all of them are not following the key, you're in deep, deep doo-doo."

Now, not all the resolvers are going to update. This is a matter of statistics. But now we've got a beautiful set of numbers. The number of users who don't use DNSSEC validation at all: Bucket A. The number of folk who use resolvers that see this sentinel and are following: Bucket B. The number of users who use this sentinel that aren't following: Bucket C. And we've got Bucket D: the number of folks still using resolvers that don't support it. We can work out the population of all those three at any time and put the unassigned bucket back into the B's and C's. We can estimate the damage for users.

Unfortunately, we can't help you fix your resolver. We don't know what your resolver does. That's your problem. From the Office of the CTO in ICANN, it's all about user damage, not fixing your car.

EBERHARD LISSE: All right. Thank you very much. When I tested it, I got legacy, and my resolvers are 8.8.8.8. and 8.8.4.4.4.4.

WARREN KUMARI: Yeah.

EBERHARD LISSE: Do you know anybody at Google?

WARREN KUMARI: Maybe. So, yes, I think they currently report that. Part of the thing is that Google doesn't do RFC5011 for key roll. [They] decided for us it's safer for us to do a manual key roll with people standing by and watching it. But there's a reasonable chance that we'll stop reporting on this, answering these questions, so that users can understand that in a more transparent way. Soon.

EBERHARD LISSE: I'm quite confident that, if and when it happens, I don't have to change resolvers.

Okay. Nick Kohli. Where are you? There you are. Hi. Nick Kohli is from Google, working on the Loon Project. He's going to talk to us about it.

And you have got half-an-hour – no, sorry. 50 minutes.

NICK KOHLI: Sure. Okay. We'll switch gears a little bit from kittens mining cryptographic keys to something a little different. Thank you for welcoming me here, Dr. Lisse. I'm sure I speak on behalf of all Tech Day. I give my condolences for your business partner. It's very unfortunate.

This is my first ICANN, so I'm quite riveted, but this is probably my 100th time to Puerto Rico, so I'm a little full of enchantment in that regard, at least. So I have that going for me.

Let's see – ah, okay. I'm, as Dr. Lisse mentioned, from Project Loon. I lead the operations group globally there, and I'm going to tell you a bit about our project and its relevance here in Puerto Rico, as we've heard throughout the day a little bit of.

The problem that we're trying to address as a project and as a part of a greater impact is something that you heard in a keynote this morning and also throughout our efforts: there are four billion people who are offline. Telecommunications has a lot to do with that. Our work in networking across the world has a lot to do with that. All of you have a great deal to do with that. If this was a global population, the person next to you would not be online, and that is a huge, huge issue. So gaining access for global citizens is something that we're very passionate about.

There's many ways to attack this problem, and they're going to all work jointly. We don't see ourselves as a one-stop solution. We'll get more into that and the strengths of different projects such as we partnered with AT&T to do in Puerto Rico. So that's a great deal of this talk.

First, I wanted to give a general overview of Loon for those of you who are maybe not familiar. Our operations are flying high-

altitude balloons with the extension equipment to be able to deploy LTE service, a cellular network. You can think of Loon as a network as a service platform that we can work with communication providers on to extend service in places and parts of the world that have nothing today.

We use a helium-filled high-altitude balloon that goes well above aircraft and weather to heights of 60,000 ft. The balloons are not stationary. They're actually mobile. One of the unique qualities of our system is that we actually extrapolate the winds to be able to take us where we want to go. So we have the ability to harness the winds, and through a lot of computing power, figure out where the winds are going to take us. We do a lot of sailing with the Internet in tow to our users on the ground.

Just in general, there's three real key components to this. One is the ability to navigate, something we spent many years on in the early parts of this project. Before we were ready to talk about this with the world, we thought, "Could we really actually make a balloon go where it needed to go when it needed to be there?"

There's a very large, complex team of software engineers – I know a few places that are good at that – and they work on this very problem, which is taking the weather forecasting and modeling today of the winds aloft in the environment and things like storms and things like global phenomena and what they do

to that to be able to ride these balloons above the jet stream in the stratosphere – 60,000 feet – and have them actually work their way up and down.

The balloons carry a balloon film. Inside that is another actual little balloon. We call it the ballonet. That balloon has a compressor attached to it. We can take on atmospheric air, thus becoming heavier, and ballast ourselves down a few hundred feet, and then release that ballast and go back up a few hundred feet.

So essentially our balloons have the ability to ride the elevator shaft of winds in degrees of hundreds or thousands of feet in an altitude range, often from the top ends of the jet stream at 55,000 feet, all the way to 70,000 feet, depending on what's optimal for us.

We're able to use that, along with the data that we get collecting wind speed and direction, along with the forecasts that are already public available for aircraft and through the weather services, to find out what winds and what streams are going to be able to take us from our launch sites, whether they're here in Puerto Rico – we do have one on the island – or a test facility in Nevada and navigate a balloon precisely within hundreds of meters to exactly where it needs to be in space and time.

Those two are very different and very complex in their own right, but we have to make them work together.

An example of that: in 2014, we took a balloon from Puerto Rico launch site and navigated across the world – I'm sorry. It was from New Zealand. We had a launch site in New Zealand and we were able navigate it right over one of our test sites in Chile within 500 meters after it had done a couple laps.

We're constantly harnessing the power of our algorithms that search out these winds and how they extrapolate that information. We're growing using a lot of machine learning and a lot of anomaly detection to figure out not only the balloons sharing the wind data they experience, but as we introduce new forecasts every six hours into the aerospace community, that forecast improves out model as well.

The second part is this clustering concept. If you think of deploying a network as a service extension of pick-your-favorite-cellular-carrier, what we're doing is bringing one balloon over our top users to be able to provide that coverage. As soon as the winds are going to drift that balloon away, another is right behind it to take its place.

That was the model for a long time for us. Through research, we said, "Well, this is simple. We just take a fleet of balloons – very expensive; a tenth or a hundredth of the cost of a satellite – and

we adopt a ring of balloons around the world at a latitude band we need coverage in."

Through all of the learning we've done, we've actually moved away from that and grown into this newer concept, which is completely turned the tides for us: we've been able to extrapolate the winds enough and go up and down enough in frequency that we can actually have the balloons create their own microscopic area that they travel in circles in.

So a balloon can find enough altitude range to cluster in the same spot, similar to station keeping, if you think of it. The balloon actually, untethered in the stratosphere, is able to go up and down enough to stay in the same relative coverage area of 40 to 60 kilometers and continue to maintain a consistent cone of coverage down on the ground. It has really been pretty groundbreaking for us in the field of high-altitude ballooning and also in the field of stratospheric or space communications.

The thing that I would leave with you that sets us very much apart from anybody that's doing anything of this type today is that, from the very first prototypes that were beer coolers of Styrofoam bought from our local 7-11 or electronics store to now to the comprehensive flight system that we've adopted and are still testing, we actually have the ability to directly connect to your handset.

We'll talk more about what we've done in Puerto Rico. The idea that you can take an existing network, an existing carrier, and expand this through the back end of the network and be able to deploy their network to their users requires no ground-based antenna to bolt onto your house. It requires no custom receiver. It requires nothing like that. It's your cell phone in your pocket.

It's mind-numbing to me to be out in some of our test sites in the desert or here and have my own phone, my own personal device, and have no service and then, thanks to partnerships with our network providers, know that a balloon is coming overhead and that I'll have full service. That's something to think about.

Just as all of us in the conference so heavily rely on mobile and data, you know it personally. You feel it. It's visceral when you don't have service. So this is something that we can all take in for our daily lives.

As I said earlier, the idea with Loon and where this concept was born was, "Where are those markets in the world that people are offline in?" There's lots of efforts in this space. The "next billion users," a term people are very familiar with, are not necessarily in geographic hubs in densely populated areas. It's a different solution for a different problem. We've really tried to break that down.

What we've designed Loon for to this date, in the frequencies and the spectrum in the communications systems that we've employed through our radios, is designed for rural and remote communities. It's designed for rural users. Whereas a smaller coverage map of towers works for highly dense areas, like downtown San Juan, what we're trying to enable is the reach for our telco partners to be able to take their network and their footprint without the heavily expensive cost of laying fiber or infrastructure and quickly deploy service to areas with low population density.

This is what we were working on for many years now, before the hurricane struck, and we were expanding at a very healthy rate. We were operating at scale, per se, already, prior to Hurricane Maria, about six months ago.

If you look at this in a clockwise fashion from the top left, there's a picture of me – I was sharing this at lunch with some colleagues – holding a latex prototype this some of those beer foam coolers stuck together as a proof of concept. That was only 2012.

As we've gone through and figured out the things we need to do for super-pressure ballooning, for the actual aeronautical parts of this, and the product evolution of our actual flight system, we were through the rapid prototyping phase in 2015. We were onto

the field operations phase, which is not exactly the glamorous software or hardware bits. These balloons, if you think about them, in the stratosphere are the size of a complete tennis court or two tennis courts wide. So they're very, very large balloons with a very large footprint.

When they go up, because gas has not expanded yet, they actually look like a jellyfish, like in that picture on the bottom left. That's not an easy thing to handle on the ground, for a variety of reasons. We don't like our skin touching the plastic because we want these things to last in the stratosphere, so we don't want to introduce chemicals and things like that. It's like treating it with kid gloves.

The joke is, when we manufacturer the balloons, our manufacturing teams actually wear shark slippers or cotton socks and stuff like that. The softer the socks, the less damage to the balloon.

Even for handling the balloons at the launch sites we've had to develop systems that are able to launch these at scale. We used to take a dozen of us and we would time it so that, when the winds were the calmest at any of our testing sites at two or three in the morning, all of us would gather around and try to launch one of these massive balloons because they were the size of the white jellyfish-looking thing at the bottom left.

That was obviously a very cumbersome process and wouldn't support our rapid deployment. So we went hard to work and actually took a giant boat crane, a gantry crane, and built an automated launch unit. This launch unit can take the same balloon and, instead of a dozen of us, take a couple operators. They can launch a dozen balloons in a day's time – in fact, one balloon every 30 minutes. So we've been able to flip that part of the problem on its head.

When we're doing that, we need to obviously manufacture more because we're building these things. We're able to grow through a lot of this scaling that we needed to do to make the timing very right for us to be able to help in situations like this.

There's not really picture on here yet – I think I have another slide talking about it – but one of the things that we were talking about at lunch is that, with the balloons, whatever goes up must come down, so we land the balloons as a regular part of our product life cycle. The balloons last 100 days or longer. In fact, we've had our record-breaker lasting 190 days. The longer we can get the balloon to stay in the stratosphere and stay aloft, the more people we can provide coverage to and the more areas we can cover. So it becomes a term of economics to figure out how long we can actually make the balloons last, and they're still a hundredth the cost of a satellite.

So we're still winning in that regard, even when we expand the day by 1, 2, or 20, and we're working balloons to last many, many, many hundreds of days longer. So that would be the goal.

I've got this example to show you and illustrate what it's like to be able to navigate the winds and explain more of that concept of sailing the stratosphere. We do have a launch site on the eastern side of this island. We've had it here since 2014, I believe, or 2015. We were constantly launching tests and we were doing with a lot of tests with a telco partners we have in Peru, learning about how the winds at the equator works. This is one of the reasons we selected Puerto Rico a long time ago.

The winds move above the jet stream, not just east to west. They move in a variety of directions. That's a novel thing that we learned, and it's a great part of extrapolating those winds. We're able to find those based on the altitudes we pick.

You can see in this example that it took us 12 days to transit from Puerto Rico to Peruvian airspace over there in the western coast of South America. Many of those days – in fact, about ten of them – we're spend just going north or south from Puerto Rico down towards Brazil, down towards the mouth of the Amazon. Then we're able to take on that atmospheric ballast and dip down into faster winds of the stratosphere and take two quick days to ride across the entire continent. Then we're able to pop

right back up – no pun intended – and are able to stick around the Peruvian airspace for 100 days.

So this is quite a standard flight. These are the ways that we're using our algorithms to extrapolate where the winds are going and what we want to do. The navigator is essentially optimizing for the fastest past there. It's not always a straight line from A to B.

There's other challenges that come with ballooning that we've had to learn and iterate over time, which is a fun part of being a technology project like this. We've instituted a whole control and monitoring mission control team. Of course, as with all things Google, we like not to always have so many humans in the loop, but there's parts of our airspace system in the global airspace system that are still heavily reliant on human intervention. So the safest thing to do is to make sure we build all the redundancies in place.

While we've automated a lot of what the flight planning and mission planning, where our balloons actually are able to sense where they're going to be in space and time and file those flight plans appropriately and dispatch to them the appropriate air traffic controllers, sector to sector or country to country – we're respecting all these sovereign air spaces – we also have a mission control team that is manning the phone, per se, to be

able to communicate in the chance of anomalies and the chance that our balloons are doing something different than what we intended or also just to be able to communicate our position for other aircraft as we ascend and descend through the airspace.

We've had to learn about the predictability of our navigation. That is continuously improving every day. I think that if I'd given this talk to you six months ago or a year ago, I would not have been able to talk about the levels of accuracy. One of the beautiful things about software-aided simulation is that we fly more miles in simulation a day than we've flown in the 27 million miles of flight that we've actually flown through the history of the project. So we're able to harness a lot of computing power towards this problem.

There's other aspects of the program that are very physical in the real world. When a balloon shows up in your farm and we've landed it there as a nice, safe space to land, working with their traffic control and the respective agencies, it doesn't look normal. It never will look normal. So it's a fun conversation that the farmer and I get to have about what it is and what it's doing.

As soon as people in the world realize that it's an access project, the reception of these things change because where we're working and where we're testing and where we've been, up until this point experimenting with this project, are in those places

that need us the most. So I can attest to you that the reception is quite warm, which is great.

And we're constantly learning. I think one of the things you'll take away from the later half of this presentation is that Loon is now not so much an experiment but open for business, which is something that, after almost half a decade, is funny to say. But we're very ready for that.

Again, to recap, before we go into the disaster response aspects of what we've done, we've been able to go from balloons that went up and pop to balloons that last hundreds-plus days on a regular basis on average. We've been able to travel the world 27-million-frequent-flyer-balloon-miles at a time. And we're delivering connectivity on a regular, daily basis, working with telco partners in regions to expand their coverage of their networks.

In the case of Puerto Rico up here until March 1st, we were working with AT&T until their network was restored to deliver connectivity right here. If you have an AT&T phone and you were on the island in parts that were unconnected or under-connected, you would have been on Loon and you wouldn't have known it. You would have been seeing your AT&T service, just as you normally can expect.

We like to say that Loon isn't a disaster response product. We never intended it or built it that way. However, one of the things we clearly learned following some horrible things that happened in Peru and in Puerto Rico is that we're incredibly resilient to the things that normally occur. We learned from earlier presentations on the ground, following disaster, the wipe-out of power and telecommunications and such.

In March, a year ago to the month, Peru had seen the worst flooding that it has seen in decades. There were landslides all over the country. These photos – a few of these photos were taken by our own team and a few were publically available – illustrate the level of how much damage there was. Entire roads and bridges were completely swept underwater. The water levels were well above roofs in many, many communities.

This was an area that we had been testing with a telco partner in Peru for a very long time, so we really felt obligated to help. But it was pretty challenging.

We had been there for a while, like I said, and we had a lot of things going for us in Peru that were aiding this effort that we took on. Our balloons were already flying in the area, and we already had some equipment that was on the ground, installed. So we were able to pipe into the main network with a backhaul ground station. That device is able to beam Internet up to the

balloons through signal, and then the balloon is able to propagate that back down to users via LTE. That interconnectivity, like I said, was already done with Telefonica in Peru.

But just as [inaudible] had suffered, we had also suffered in terms of what had happened on the ground with the weather. Our team was able to take that situation, do some survey work, quickly get our backhaul sites up, and, through a lot of work going on locally in Peru and our Mountain View headquarters, get the ground station and backhaul connectivity restored so that the network that was already flying above the Peruvians was able to be just connected again, turned right on. So we were very successful.

It clued us into laying the groundwork, the first bricks, for realizing that we were able to help when a disaster hit a little closer to home for us. Like I said, we had a launch site here in Puerto Rico for a long time. When Irma and Maria hit in Puerto Rico, our site was genuinely affected, devastated to the same degree. It was a brutal storm.

These are pictures our team took. There's a picture of the red Jeep that you can see was a drive to our launch site. It's not every day that a boat is sitting in the middle of your commute. So it really provided some reality checks for us on what our team

was going through on the ground and how bad the situation was to be without power, drinking water, and that connectivity.

We saw this earlier, too, so I want to reinstate that a little bit – that message.

This is where we mention Status.pr. There's a different illustration of that, not the satellite version but the screenshot that they were actually providing. Really, at this point we were starving for information, any time in a blackout situation. Just south of our site, in Humacao, you can actually see some of the ground-based equipment on the ground. These are towers completely knocked over and telephone poles split in half. So, again, it's some of the devastation.

We were on the ground six days after the hurricane, after our folks had left their shelters that were here that stayed on the island and the folks that evacuated. Like is said, we were back six days after the storm had passed. We had a lot of things going for us.

One of the things we had going for us was that we have a prototyping launch site, sort of a test facility, in northern Nevada. That site isn't usually used to reach our networking partnerships that are going on around the world because we have this site here in Puerto Rico. So the proximity is always a little closer.

However, since our launch site was down, we were able to quickly turn our launch site in Nevada into the production facility to be able to support deploying balloons over the top of Puerto Rico, something we had never really considered before.

Because we had people here, they were able to pull our equipment out of the bunkers and storage and quickly engage it for a while, using satellite backhaul and quickly working with AT&T and T-Mobile and other partners, including the GSMA, the FCC, and the Puerto Rican government to find backhaul access because the undersea lines were not affected.

Like I said, within a matter of a couple weeks, especially with the interop with our telco partner, we had balloons covering the island. We had service restored to users, focusing quickly on the critical communications piece and then opening up to, like I said, AT&T and T-Mobile to be able to supply coverage for their regular users.

We have a pretty mission-oriented team. They're based in Puerto Rico. They shared some insights with me. I asked them for some cool stories, and I remembered these three as ones that stuck out in my mind to try to illustrate the resolve of what our project has been through.

We had a team that needed to get telecommunications gear in. Of course, we hitchhiked it onto boxes of water filters that were

being sent to San Juan. We joke about the most expensive water filters and radio equipment that's ever been sent into the field.

The concept of just simply getting radio equipment for our ground station, a replaced part that might be the size of your laptop, into an island that has no working airport, that has no power, that has no communications, that has really nothing going for it is incredibly hard. Finding unique ways to be able to get that back on – I can't imagine, if I had that much trouble, a multi-day effort to get a shoebox-sized piece of equipment following a hurricane, how much trouble the power companies or the network companies had to get cell towers or other replacement equipment back here. So we were very fortunate that we in a position to help with that.

We were partnering with the military, who were working alongside our launch site over there at the Roosevelt Roads decommissioned air station. They reactivated that air field, since that runway was unaffected.

Here's a couple of pictures. We had to do some trade and barter. We really needed some support from them to be able to get into the airfield and get our power and lights back on so we could start figuring out our ground station restoration. They needed a crane to be able to help one of their rescue helicopters. Well, we

had a crane that launched balloons, so why not attach a helicopter component to it and use it to do such?

So our team was pretty resilient in that matter. All of that shaved time off our response. Like I said, there was a lot of different work that went into creating the network that has 200,000 users on it on an active basis.

Here's a map of our coverage in Peru on any given day. I think I took this sometime in January after the response on a random day, just to illustrate how that coverage looks like for those of you learning about Loon. The balloons represent a cone. In that cone, there's a [confidence] to the edge of that green circle. If you're underneath that, that's where you're going to get LTE penetration.

I have a short, two-minute video that we made to help illustrate some of the efforts we had in Puerto Rico that I thought I'd share before we take questions.

UNIDENTIFIED MALE: When I really realized what was going on, the next day I was trying to call my dad. I was trying really hard to communicate with him, and I couldn't. This was me, a government official, who was not able to communicate with his father.

UNIDENTIFIED FEMALE: I didn't expect that it was going to be this bad. I knew it was going to be bad, but I wasn't expecting that it was going to be this bad.

We had Internet before Maria. We had so many things that we depended on the Internet for everything.

UNIDENTIFIED MALE: We actually have two launch sites. One is in Puerto Rico, and one is in Winnemucca. What happened when Hurricane Maria blew through is that it took out a substantial portion of our facilities in Puerto Rico. So we switched to launching out of Winnemucca, and we had to find new ways to navigate the balloons down to Puerto Rico to respond to the emergency.

The balloon actually wasn't designed for disaster response, but because we already had ground infrastructure in Puerto Rico, we were able to quickly respond.

I'm excited that were able to provide basic connectivity to hundreds of thousands of people after the hurricane.

Some of the things that we learned from this experience are that, before a disaster strikes, we need to secure spectrum, import all of our equipment, and test the full system so that we can restore connectivity as quickly as possible.

UNIDENITIFIED MALE: There's got to be a playbook for disaster. Telecommunication has become a necessity. I'm a firm believe that Loon should be part of that disaster response playbook for telecommunications.

UNIDENTIFIED FEMALE: We want to give back. We have to give back. We were doing it before, and we can do it again. The Internet is going to do that.

NICK KHOLI: So that's what I had to share. I'm happy to take questions if folks have them.

EBERHARD LISSE: Thank you very much. We'll take it in order of –

UNIDENTIFIED MALE: The goal of Project Loon –

NICK KHOLI: Oh. Sorry.

UNIDENTIFIED MALE: Sorry. I've heard rumors over the year of shortages of helium. Is that an issue for you guys?

NICK KHOLI: We were actually discussing this at lunch. We do use helium gas. We've certainly looked into all types of gas, but the amount of helium that we use compared to the actual global supply is quite negligible in this regard.

Things like being able to figure out and optimize how many balloons were need – we're not talking about tens of thousands of balloons in a ring. We're talking about using the winds to cluster a few balloons – in this case, five – to supply a whole island. So it's quite small in that regard.

We're constantly looking at that situation, too. There's all sorts of developments in that space, and we know it's a finite resource.

ARTEM GAVRICHENKOV: Artem Gavrichenkov from Qrator Labs –

EBERARD LISSE: Can we please identify each other or ourselves before we speak? Because the remote audience doesn't know who's speaking.

ARTEM GAVRICHENKOV: That's what I had just done. So repeating: Artem Gavrichenkov from Qrator Labs. I wanted to ask you to please upload your

presentation to the website. It's very interesting, but I can't find it. So please do.

NICK KHOLI: Sure. I think Dr. Lisse had mentioned this in the morning. I apologize. One of the caveats to coming here to talk about this is that so much of Loon is changing all the time. I could point you to our website, ProjectLoon.com, but the team at home is actually quite sensitive about what they share. They were happy to have me come and share this here. At the end, I can share a link. I'll make sure that goes onto the Adobe Live to be able to point you to our website where you can find more info.

ARTEM GAVRICHENKOV: Thank you.

ROBERT MARTIN-LEGENE: This is Robert Martin-Legene from PCH. You have one balloon covering a radius of 40 kilometers. I'm wondering: in this case, do you cover it with all operators, or do you work as a repeater? How does it work? How does power consumption work? I suppose you have solar cells and batteries for the night. How many users can you actually support? Because one balloon and 20,000 users by five balloons...isn't a little bit oversaturated with everybody doing weird stuff and –

NICK KHOLI: I'll take those questions one at a time. The first one was regarding square coverage, correct? In the model that we have, the radios are attenuated to be able to support reaching direct to the handset. That's where the footprint changes. That rural focus we have is based on density, as you've caught onto. The density problem is something we constantly are able to manipulate, based on what the optimal coverage in regions is. It's taking a population map and figuring out where users are, working with a telco provider.

Actually, to go to the very first question you asked, we dedicate one balloon per network. In the case where we working with T-Mobile and AT&T, we might have a balloon supplying through their core network and then another balloon supplying through another telco operator's core network to keep the data separated. So it's not a completely shared pipeline because it has to interconnect through the operator who is responsible for the provisioning of your device and your data and all of the flow that goes on behind that.

In essence, there's nothing says we couldn't put two balloons side by side and one would serve AT&T customers and one would serve T-Mobile customers. The penetration piece is something that we have a good sense on with the operator to

collaborate with them and find out where their users and what we need to support. Even in the mobile market, I feel quite confident in that space and being able to understand it.

The second one: the reason why it's not a beer cooler any more is because we need to take on enough battery and radio optimization technology in the flight system that you saw that's not a prototype anymore in the video. With four solar panels and being quite a large piece of equipment, it's still much smaller than a cell tower or a standard ground-based infrastructure. But it's quite enough to be able to support us, providing through the night and harvesting power from the sun during the day.

One key advantage to being 20 kilometers up is that you're a little bit closer to the sun, so you can get quite a bit more energy than you would from the same-sized panel and the same grid on the ground. So we're able to harvest in that way.

Sorry – did you have one more?

ROBERT MARTIN-LEGENE: No. I was thinking about how most of the island would have been relying on most of these balloons. Weren't they completely saturated with traffic?

NICK KHOLI: Yeah. Well, for one, there's multiple ways to think about how that saturation happens. You can have it in connectivity in disaster. You can have it to a select group of users if you needed it to be for emergency response in some instances. We talk a lot now – the narrative has shifted – about preparedness. I think Dr. Lisse was mentioning that, too. That's something that we think about that we're making conscious about.

Two, it's really about the diameter of the pipe, to say where your backhaul is coming from and the latency and speed of that backhaul. When we're able to provide that, we know what our saturation numbers are. Then we're able to figure out whether need two balloons to provide coverage for this area, or three, or ten, or five.

So it's an algorithmic problem. You plug these things into this computation and say, "Oh, okay. Based on the density and the type of traffic being consumed, I know how many balloons I need to serve for that many users."

Again, the experience is different than the data-rich experience that you might have in the room right now. We're using LTE spectrum and LTE radios, but, really, the user experience is probably a lot less speed than that. You're probably thinking about a 2G/3G comparison in terms of how traffic is flowing. But it's connectivity. It's the critical messaging. We will support the

top messaging services. We'll support e-mail. We'll support things like that.

ROBERT MARTIN-LEGENE: Right. Okay.

NICK KHOLI: Thank you.

ROBERT MARTIN-LEGENE: So one balloon may be, like, half-a-million users if you're lucky? Is that something you can estimate?

NICK KHOLI: I don't think it's something you could just estimate – a binary like that. It really depends on how much traffic is flowing and what your backhaul situation is like.

HOWARD BENN: Howard Benn, Samsung Electronics. First of all, I think I'd make a comment about your numbers. At the moment, mobile connects 6.2 billion active users. So I'm not quite sure how many additional users, because you're using the LTE technology – something like half of the users don't actually use LTE today.

NICK KHOLI: Yeah.

HOWARD BENN: But more importantly, for 5G at the moment, we are developing something called NTN (Non Terrestrial Networks). So we're trying to optimize 5G to make sure that, when that technology gets into handsets, it will cope with situations like this.

I was wondering why companies like yourselves getting involved in setting those standards, given all the experience you've got with LTE.

NICK KHOLI: I think we've done a couple of things. I'm happy to discuss more offline and connect you with those specific people on our team. I know that we're heavily engaged with the GSMA and with other bodies on setting those standards. But we also are very reliant on some of our partnerships in this space – the radios and things, working with Nokia. Our team just came back from Mobile World Congress. I know that they're heavily involved in that space.

What the future is? The beauty of – again, going back to a system that goes up 100 days at a time – it is that we're able to refresh the hardware just as fast to keep up with these changes. So it's not just things like software we're pushing. It's very common

that, by the time a fleet of balloons might land, we're pushing hardware changes into their radios. We've seen that iteration over the years as we've developed new radios.

The first iterations of Loon, before we pivoted into the LTE space to go to handset, were actually with Wi-Fi. That's why it's such a novel switch for us. We originally started with the old Ubiquiti rocket and tried to shove that in a box.

So we know the space is changing. I'd like to think that we have that covered. I'm happy to connect offline. We can [pass people] that.

UNIDENTIFIED MALE: [inaudible]. Two or three questions. One, how do you cover under the licenses if there is any set license in a particular country to operate [foreign interest] services? Does this require any license or do this bypass any license? Or can you just fly anywhere and provide services? That's one.

Secondly, how does –

NICK KHOLI: I could take them one at a time if it's easier for you. One of the funnest things about my career at Loon is that we're at the intersection at many of the world's most heavily-regulated

industries. I'm just kidding – it's not that fun. But the point is that aviation telecommunications and broadcasting and all of these spaces are quite heavily regulated. So we recognize that. We also recognize that we span across them.

On the aviation side, we are beholden to the sovereign airspace controls of every country we fly over. Hence, we don't fly over certain countries.

In cases of natural disaster, to some of our countries that normally tell us to go away, we're certainly open to saying, "Hey, Puerto Rico needs you. Despite foreign relations, let's make this work right now," which we're pleased with and proud of; that people can recognize that in a disaster scenario.

But in terms of being a product, as an actual market solution in the enterprise space here, we have our wind algorithms that block off those airspaces and are able to react to the daily change of governments and issues around the world so that the balloons don't have to overfly Russia or North Korea or China or those places when those things are not allowed. That's the beauty of being able to navigate.

On the telecommunications side, switching gears, we secured these licenses with the FCC – and experimental license to transmit. We worked with FEMA as they were using some of the

available unlicensed space for short radio links on the ground. So we are able to deconflict a lot of those things.

We do have to get those permissions. A lot of the spectrum that we secure is through our telco partners that we work with because they already own that spectrum for their users. They're using it on the ground. So a lot of time we spend, engineering-wise, is in deconflicting how those signals correlate with each other or work atop of each other.

UNIDENTIFIED MALE: Thank you. Second, I can understand the case of disaster recovery, where Puerto Rico and Peru faced the situations that you solved. Do you think there are very difficult situations and locations where the fiber has not reached, even the wireless hasn't, and even the fiber has not reached? Have you ever served places like that?

Do think it is a viable proposition in terms of the cost and everything else to sell them [inaudible]?

NICK KHOLI: Yes. Absolutely. The first place we were going with this, before disaster struck, was working with these telco partners to illustrate how we can reliably bring continuity of services that they expect you demand a degree of reliability working with

your regular market customers in places that they actually don't go.

The reason people aren't online, say, in a valley in the mountains is because the cost of laying fiber to get to that limited number of users. If we can democratize that by making it a little bit cheaper for the telco partner to be able to continue to serve those users on demand with network as a service in this concept, you can go places that you have to take towers and power and cable to and still be able to get – those people have handsets. They still have an LTE device. In Peru, we learned that there's a large variety of our customers from our Telco partner that live high in the Andes and come down to the towns on the coastal range in Peru to work. They take their device with them in their pocket.

So they are used to, in the third world and developing countries in the second world as well, having the devices and knowing when and where they can't get service. Anybody from a rural community knows the spot you can go to and the weird stance you have to be in with a raised leg. I think there was a picture of it one of my very first slides – of finding that tree so that you can get that signal. That's a problem that we think we can solve.

UNIDENTIFIED MALE: The last question: I can understand that the balloon will be having two transmitters and two receivers to [interact as an intersect]. Can you connection any of the transponders and receivers in any geography in the coverage of this thing?

NICK KHOLI: Currently, right now, in LTE they talk about it in phases. We certainly fly one transmit and one receive. We've looked at multiple, up to four sectors. So we're constantly evolving that technology. We're experimenting with all different types. Flying multi-sector antennas is something we're very excited about.

Again, it points right back to that, if this was a $100 million satellite or something in this space, one of the things that makes us so competitive, we feel, is that we're able to refresh this hardware and figure that out. So there's so much iteration going on. So, yes.

UNIDENTIFIED MALE: Okay. Thank you very much.

NICK KHOLI: No problem.

EBERARD LISSE: I've got two comments. A friend of mine from Namibia came to Puerto Rico and to St. Martin. He works for a company that uses TV white space. He told me that regulation is not a problem. This is not a developing country or India or Pakistan or Namibia, where the regular needs to be [stroked] and sorted out. In emergencies, things in America tend to happen very quickly – at least the government officials, not necessarily the politicians. But regulation is not a problem.

UNIDENTIFIED MALE: [inaudible]

EBERHARD LISSE: Hang on. The second thing: I fully agree. When I go and visit my mother-in-law in rural northern Namibia, I'm always watching – I can't do YouTube. I can't do this. I can't do that. She doesn't care. She can phone. She can send me a telegram message – we don't use [inaudible] that works even with 2G.

We're probably a little bit spoiled with the requirements we have. I remember [inaudible] PC and we didn't even know what an attachment was. Nowadays, there are 10 or 20 megabyte-attachments on a regular basis. No drama.

As you said, it's visceral. If you go somewhere and you don't get 4G – it's instant gratification. Maybe if we come down a little to

earth and provide a service that is available to everybody if it's not 5G or 7G – it's just 2G or 3G – you can make your phone call. If there is an emergency, maybe you don't an ambulance. But a helicopter will usually find its way.

These things work. As for the rest of the presentation I can just say, "Wow."

NICK KHOLI: One of the biggest motivations our team in general had for trying to deploy so quickly in Puerto Rico was that Status.pr site. If anybody has navigated that site or if you could do it right now – the presenter said it was up; Mr. Rodriguez – the site is super plain. It doesn't take a lot to load. It's not fancy. It's primarily text-based. But it told you where the water was. For Puerto Ricans suffering from this disaster, they weren't trying to YouTube or FaceTime, although we support those services. We support data. They were trying to find out where the gas was available or the water was available for their generator or for their home. Or which hospitals were open.

So I think we feel really strongly about knowing that some access is good access.

EBERHARD LISSE: Thank you very much. This was one of the presentations I have ever seen here. Thank you so much for coming.

NICK KHOLI: Thank you.

EBERHARD LISSE: Give him a big hand.

We are now going to allow the other people to leave who were polite enough not to open the doors all the time while somebody was speaking. We meet at 25 past three.

[END OF TRANSCRIPTION]