PANAMA – Tech Day (2 of 3)
Monday, June 25, 2018 – 15:15 to 16:45 EST
ICANN62 | Panama City, Panama

ROD: OK, I am Rod [inaudible], I don't know who that [inaudible] guy is he sounds pretty rad though. SSAC chair, and I've got [inaudible], and Warren [inaudible], and [inaudible] up here. We're going to talk about a really cool attack, cool from some perspective, as Warren said, there's nothing new here, but it's like well yes, somebody proved that it could be done a long time ago, somebody actually used BGP hijack to create, to redirect queries to Amazon's DNS service, to then put up a DNS server to redirect for a particular domain so they could steal stuff. That I consider pretty cool. This is our emerging security threat, and we'll see if it's emerging or not, if other people do this going forward but it's one that touches on a lot of things that have been big problems for a long time so we thought we'd talk about it. One of the things SSAC wants to do, is bring more security topics to ICANN meetings, this is a first shot at that and we'll see how this goes. We may expand the program going forward, so, can I get the next slide. Have we got the slides or does... yeah... cool, alright, I've got control. So, this is the standard... is there anybody in the room who doesn't know who SSAC is, what we do, and all that? Do I need to explain it, raise your hand. Oh

yeah, if you're members you don't get to do that. Alright, so good, I can skip this slide. We do have 101 publications now though, so we've broken the century mark. Actually we broken the century mark, I think just at the end last year or the beginning of this year. Is that where I introduce people, and we're going to talk about these things and for those of you, we're going to give some background around BGP and BGP hijacking, if you are not familiar with that or you are and it's a refresher, and talk about the attack and then various things you can do to detect and potentially mitigate some of this stuff then overall help them. We're going to do a panel at the end of, with questions and discussions about the various issues. So, OK, I already did that... Danny was supposed to be here but I think he's got a NomCom vote or something he's got to do, so Danny McPherson, if you see him tell him to get up here. So, he took part in this. We drafted Warren at the last... I think he showed up on the phone call by mistake and he got drafted as a result.

Anyhow, let's talk about the highlights here, so these are a couple of the articles that came out, there was a couple of really good ones, talking about this attack and actually pretty good coverage describing what happened. Got a lot of play on the DNS operations list and a few other secret score lists, etc, but the net effect of this was $170,000 I believe in Etherium was stolen in a matter of a few minutes, by this... you could call it a

phishing site that was set up, basically a copy of the real site but it wasn't your standard phishing attack. That's what they did, they created that fixed site, they stood up this... so it's a cryptocurrency, I'm not sure if you've all heard of Etherium, but it's very much like Bitcoin, 170k there you go. But they never touched the actual site, they never touched the actual infrastructure of anybody involved in the site, they just did the hijack. So it's something that you're just not set up for from a defensive posture very well, and the end users don't necessarily, there were some warnings, but they were the kind you click through, and we know how people are good at just clicking through little warnings that come up. So, it was really hard to detect and do things about. So, that's the overview, we'll dig in a little bit more into the details of how the actually pulled it off, but at the end of the day they were able to get into the wallets, literally, of these customers who were just logging in as they normally would, and were able to siphon out their coins by intercepting those logins and replaying those on the real site and transferring the Etherium. A very cool way to make a lot of money real quick if you know how to play with the internet infrastructure. So, I'm going to turn over to Warren to talk about BGP, [inaudible], OK.

UNKNOWN SPEAKER: So, what was really interesting about this particular attack, is that it was the routing hijack that really initiated being able to do a cache poisoning attack. So, typically people look at, OK, here's a BGP issue, a routing issue, here's a DNS domain issue, but really recognizing that there's a lot of correlation between these different networking, fundamental protocols and understanding where you can make a difference, right? In your realms of authority. We've had the DNSSEC workshop, for I don't even know for how many years here at ICANN, people have been talking about RPKI, we;ve been talking about BGP filtering, there's all kinds of best practices you can do for routing and for DNS, and none of it's a silver bullet. So, what this particular attack showed, is that you do have to pay attention to the best practices for the routing infrastructure as well as the DNS infrastructure to really protect yourself from these kinds of attacks. So, for those of you who may not be routing experts, there's a routing protocol that was created in the early '90s called the board of gateway protocol, and it's become ubiquitous to inter connect the network of networks that make up the internet. So, each network is identified by any anonymous system number, AS number. You can think of it as, if you have an enterprise and I'm over simplifying greatly, each enterprise would have an AS number. They have to be unique and as a funny story, I used to work at Cisco Systems and you used to have examples of AS109. So, a lot of network enterprises

has AS109, because they didn't realize that they had to go and get one, and of course, then routing broke all over the place. But, an anonymous system number, uniquely identifies an entity that is then, gets a certain IP address range, a cyber block and then is allowed to route it. Now, is allowed to route it, I am going to put that in air quotes and I will get to that in a minute. But, each AS asserts the reachability for the destination to which it provides con activity, and the big assumption here is that you're allowed to assert that I can reach these networks, or I own these networks and you know, come through me to get to this particular cyber block. There is no central authority or point of control, and that's a really important distinction to make. The regional inter registries, you have your LATNIC, your APNIC, your AFRINIC, RIPE, and [inaudible] allocate the IP address blocks but they do not have any kind of operational role. So, again very much over simplified, right, but if there is an AS that has con activity to a slash 24, their 19202.0 on the left, it will then asset that hey, you know, I'm the one that can reach this particular network and would then advertise that particular route to it's connected routers. Then the same for the other ones, so for example AS 64501, it is inter connected to the 19851.100.0, so it can reach that network and say yeah, here, advertize that route. It also is one hop away from the 19202.0, so will also advertise that it can reach that particular network.

So, again very much oversimplified, most enterprises have a lot of different addresses ranges that they have reachability for, or rather ISPs, sorry I'm mispeaking here, but how BGP works, it will just say I can reach these particular networks and then it will give some kind of a way to the route, where it will say I am directly connected and directly connected usually are more believable than hey, you know, I can reach it but via this particular other router or path. So, what exactly is a BGP hijack? So, both routing and the domain name system protocols seem really simple, right? For BGP, hey, I'm just trying to figure out we're unreachable. I announce a route, I see that I can get to certain destinations via this path, simple, right? But, it's the same thing as saying, OK, DNS, well hey really simple right? I know a name, I get an IP address, and I know where to send the packet to. When you start looking at the intricacies of both protocols, it's quite complex. So with BGP there's a series of mechanisms where you determine what the best path is, but for this particular example what we're looking at, is that the longest prefix match is always the one that's preferred. So, for example, there's also something called [inaudible], so with routers, right, you don't want to announce all specifics, because your routing tables will get so large and the memory in the router won't be able to hold all the routes, so best practices is to consolidate and aggregate the routes before you announce them, and from the policy is that a slash 24 is the longer prefix that actually any

route will accept. But, a lot of routers, or a lot of enterprises or ISPs, they will either announce a slash 22, slash 23, slash 20, whatever have you, to conserve the routing table. This makes it possible to actually insert a longest prefix where somebody can hijack where they can then announce a slash 24, even though the entity that is allowed to send, or that's responsible for certain address range will try to do the right thing for the internet and announce a slash 20, or a slash 23. So, a BGP hijack occurs when an illegitimate route is advertised and preferred over a legitimate one. Often, this happens because people, operators make mistakes, fat finger, like you just type the wrong address in by accident, and this happens more often than you would think. Usually, people catch it fairly quickly, there have been examples in the past, like a Youtube incident years ago by Pakistan where they fat fingered and you blackholed Youtube, nobody could reach it. Again, that was a mistake. What this particular incident showed, is that people are starting to do this deliberately to cause malicious activity, and this is where a lot of eyebrows were raised to say, this can be a much bigger problem in the future and which is why us in SSAC wanted to raise thIs to the community, to say you really do have to start paying attention to these BGP hijacks, how they can affect the stability of the DNS and what can we as a community do about this?

So, we look specifically at what happened here at this Amazon route 53 attack. So, you had a number of Amazon route prefixes that were hijacked, so Amazon was announcing slash 23's and somebody deliberately did a BGP hijack where they announced slash 24's which meant that they now had control of the routes, and what they ended up doing is that out of that slash 24, they then created fake authoritative DNS servers that then did a DNS cache poison to the recursive resolvers so that when an entity wanted to go to Etherium, they used the cash poisoned answer, and would actually go to the authoritative server that was fake. That allowed for the malicious activity to take place. So, if we look at it from a picture perspective, a picture always tells us a thousand words, so Amazon usually would announce a slash 23's, and that's shown on the left, so on the right, there's a malicious actor that then interjected more specific routes, slash 24's, that would then get accepted by the different routers. Now, is it appropriate that they would accept it? No, and one of the things Warren will talk about is one of the mitigations is that people should be doing prefix filtering on the internet. People are horrible at doing filtering. I used to do a lot of security workshops, years ago, and I think 10 years ago, I remember, I was so frustrated by people not filtering that I would talk about filter, what are we going to do filter, wait, what are you going to do filter. Because, I wanted to get people to really start thinking about filtering. They got so annoyed at me that one man raised

his hand and said mam, if I promise to filter will you stop talking like that? I said, yes, but now I want to see your configuration tomorrow. To continue, right, if you're not filtering, you keep propagating the bad route so then you've got everybody believing this hijacked route and this is how you can then do other subversive activities, in this case creating DNS cache poisoning. What ended up happening is that the victim would then want to go to Etherium site, it would then send a query out, how do I get to Etherium site, the recursive resolver would also then go hey, how do I get to it? It would be routed to the malicious authoritative route 53 DNS servers, which would then do a cash poisoning, send the fake answer, and the victim is going to the malicious site. That's it from me. Warren? Sorry, one more slide. One of the things to really recognize is that while a lot of people talk about cryptocurrencies, it's good, it's bad, I trust it, I don't trust it. I don't care about that. What to me, it showed that this is a hijack that can cause cash poisoning for a numerous amount of issues and you can do farming attacks, email interception, credential theft, what have you. What I worry about is that if we don't do the basic network hygiene for either routing or DNS, that we can be in a world of hurt. That's worse than just this one Etherium issue.

WARREN: Now I've got the clicky thing. So, if people have been paying attention to this, hopefully you are all now terrified. But I am here to tell you it's not actually all doom and gloom, it's only almost all doom and gloom. So, in order to actually figure out or to be able to mitigate the problem, you first need to know there is some sort of problem, so all of the BGP data is public, in order for your router to be able to route packets somewhere it needs to actually get the BGP tables. This means that you can look at the tables yourself and figure out if somebody is announcing your address space, but that gets kind of annoying, and tricky to do. There are a lot of services that are set up specifically to do BGP hijack monitoring and they will happily send you alerts and things like that. Unfortunately, often for a cost, but you know this is probably worth spending some money on. Here's a list of some of them, one of the ones which I think is funniest, is basically a stream of BGP information sent over Twitter. People have made this as simple as possible for you to be able to do your own analysis and monitoring work. There are a bunch of other services as well, thousand eyes is a fairly well known one, BGP mon is one, and also a number of the RIRs, the people who actually hand out blocks of addresses, for example, RIPE have their own services that you can subscribe to that will let you monitor for BGP hijacks and things like that. Also, obviously if your route gets hijacked and all of your traffic gets steered away somewhere else, one of the things that you'll notice, hopefully

fairly quickly, is the fact that you're no longer getting the traffic that you're expecting to get. All of the traffic that you're getting suddenly drops to zero, this is probably a bad sign, in fact this is definitely a bad sign, possibly evidence of some sort of route hijack something you need to follow up with, or a [inaudible]. Then, another obvious thing you can do, is just monitor your own DNS infrastructure from all over the world, you can do queries [inaudible] is a very well known service. You can ask [inaudible] to go off and resolve a set of names for you. A, you can see where those queries flow to, if they actually hit your name server and that is probably what you're hoping for, but B, you can also make sure that the answers that RIPE is getting are the answers that you were expecting to be getting. This is probably something that you should be doing regardless of if you're watching for BGP hijacks.

So, now that you know that you've actually got a problem, what are some sort of things you could be doing to deal with it. Actually, that isn't what you should be asking yourself at all, the question is what should you have been doing before you discovered the problem, and what that is, is things like BGP prefix filtering. If you are an ISP, or provide services to someone else, you should know what addresses they are going to be sending you, what routes they are going to be announcing, and you should only be accepting the routes that you should be

getting from that customer. In the route 53 hijack, if the upstream ISP of the attacker had been paying attention, they would know that that customer should never be announcing Amazon's space, so should never have accepted the route. As with many sort of health and safety things, this is something people don't bother following up on. Everybody I am sure flosses and brushes their teeth for two minutes everyday, whereas I am sure everybody always applies [inaudible] filters. It's one of the things that people mean to do, don't necessarily spend as much time and effort as they should. Of course, this causes problems for other people, unlike with other brushing your teeth only causes your teeth to fall out it causes problems for other people as well. So, ISOC and a bunch of other operators have gotten together and organized something called MANRS, mutually agreed norms for routing security. Basically, what this is, is some efforts to agree on what the correct set of MANRS are, if you're working on the internet. This includes a bunch of things like filtering your [inaudible] correctly, making sure you've got reachable abuse contacts and things like that, sort of generally what you should be doing to actually be a good citizen on the internet. A fair bit of the MANRS stuff talks about BGP prefix filtering, and that's because that's one of the problems which has been around largely since BGP existed. As [inaudible] said, some of these are malicious, the huge majority of prefix hijacking, or leaking of filters are purely accidents, but you

know, what can be done by accident can be done maliciously. It was about URPF, which is somewhat related to route filtering, except that that actually says you should only accept packets from users who have certain address ranges, not just the routes themselves. Another thing is the RPKI, resource public key infrastructure, which we'll talk about in a bit more, and then another mitigation technique which is probably not necessarily a best common practice but which everybody does anyway, is to use the longest prefix possible for critical infrastructure. So, when people do prefix hijacks, often what they will do to make sure that all the traffic flows to them, is that they will announce a longer prefix. EGP, what is routes according, actually all of IP routes, routes according to the longest or most specific prefix, and so if you make sure as the legitimate service owner you're announcing the longest possible prefix, people won't be able to hijack your route. However, that's being slightly anti-social, because the more prefixes people announce, the more space it takes up in everybody else's routing tables. This is something that everybody, or many people kind of do, but isn't necessarily the most social thing.

Mitigation techniques. A fairly obvious thing is if you spread your list, or your DNS recursive servers across a bunch of different routing announcements, people will have to hijack a routing announcements, to be able to take over all of your traffic. So,

you know, have your DNS servers spread in a few different AS's. This is actually something that is recommended practice, regardless of this, because if you have a network failure, it's unlikely to take down multiple. Very similar, use multiple providers, do your own DNS and also have an external provider, and also server related to a previous slide, just keep watching to ensure that the amount of traffic that you're getting hasn't suddenly dropped, and if it has, you probably want to figure out why. For the [inaudible] wallet, a really good mitigation would have been, if the name was DNSSEC side, the route hijack redirected all of the queries to an attackers DNS server, if there has been DNSSEC in place, anybody who is doing validation would have gone to the malicious server, would have tried to resolve the name and would have seen the DNSSEC signature didn't match and would have no report at that point. Obviously this requires you to sign and people to validate. Another thing is, the [inaudible] site was signed... sorry not signed... had a HGPS certificate, was a secure website, the attackers didn't bother putting up a HGPS cert, probably it would have been a little bit tricky for them to get one. But, unfortunately, users when they see a big red pop up saying this isn't the site you want to get to will just click through it. This particular case, they didn't even get a pop up, it was just an insecure site. Something that could have helped in this case, is if [inaudible] had a [inaudible], which is basically sort of almost like certificate pinning, we use HGPS

and we will always use HGPS, would have been a good mitigation. Users would then at least, all users would have seen a, this is probably not the site you want. Related to DNSSEC if DANE was in wider use the attacker wouldn't have been able to actually, they would have been able to move the site somewhere else but nobody would have been able to trust it. Also in the motherhood and apple pie, all ISPs really should be filtering your customers, and if you are a large-ish DNS provider and speak BGP, you should really be filtering the routes that you get from your provider. So as [inaudible] was saying earlier, filtering is something that everybody says that they do, or at least almost everybody says that they do, unfortunately as we've seen by the fact that route hijacks continue to work, not all ISPs filter, or filter as well as they should. They don't actually filter between providers, a lot of them don't filter their customers, or they have sort of special one off cases where they filter all of their customers, except for those few weird people who are kind of on the side, and then that serve group gets larger, that filter very reliably.

Next slide has sort of set on where you really should be filtering, so ISPs should filter from their customers, customers should filter from their ISPs. ISPs should filter to their peers, peers should filter from their ISPs. This is how the world should work, unfortunately, in the real world it doesn't always look like that,

things get a little bit messy. Part of that is it's often hard to actually tell who you should be accepting routes from, so if you're an ISP and a customer comes along and says I have this address space, it's a little tricky to know whether they really do have that address space, when they have got the address space from the RIR, and then that got given to a friend of theirs, who then lent it to them, and they said you can use it because I am not. It gets really complex. There are various ways that this currently gets authenticated, letters of authorization is one, where the customer sort of signs something saying, no really I am allowed to use this address space. Obviously, if the person is malicious, they are going to be perfectly happy to say yeah, sure I am allowed to use Amazon's address space, whether they are or not. There are some address spaces or blocks of addresses which we know are not currently being used and won't be used for the foreseeable future, for example on the public internet, the 10 net is not something that will be announced, or should never be announced on the public internet, 192.168 is sort of another similar and well aware of. There are various other blocks that are used for infrastructure which should never be announced. There are some templates here on a set of the filters which you can just apply and know that they will be OK, the [inaudible] top one, is sort of the best known filter. Oh yeah, a term that you happily use without using. A bogon is basically just a set of prefixes which you know will never show up on the

internet, I don't know where the term actually came from, except everybody has always called them bogon's. I think it came out of a Science Fiction book, Star Trek, a little sad...

UNKNOWN SPEAKER: Warren, Google is your friend.

WARREN: So, let me go back one. So, when everybody... whenever somebody gets an address block there is a thing called a routing internet registry, where you should go along and register it, basically IANA gives blocks of address space to regional internet registries, they publish these in IRRs. When an ISP gets one it gets published, if they give it to a customer, that gets published. These are not always very well maintained, there are also a bunch of different ones, so you will find a route showing up or a prefix showing up in one of these IRRs, saying a prefix [inaudible]. Then in a different one, it claims that the same prefix belongs to Fred. It's very hard to tell which is indeed the correct set of information to listen to. Many of the RIRs also require that you give them a substantial amount of money to register your one little prefix, like $2,500 is an example, and if you've only got one or two prefixes, it's a lot of money to put... pay $2,500 a year to say, hey this is really my address space. They also, the syntax that one uses to talk to the IRRs is a little annoying, and so a

bunch of folk came up with something called RPKI. This is basically taking the same ideas that you use in the web, the sort of PKI around HGPS and applying it to routing information. An actual better sort of analogy is that RPKI is kind of like DNSSEC but for the address space. So, with DNSSEC the route is signed and then dot com is signed, under dot com, food.com is signed, etc, and each layer sort of signs the layer below it, or at least signs the fact that the key is valid for the layer below. Routing people are going to get all kinds of twitchy with this analogy but it's a fair amount of similarities, the IANA gives a bunch of address space to various RIRs, the the RIRs take chunks of the address space, give it to their customers who are often ISPs. Customers take parts of that space, give it to their customers. So you can sort of build something that looks almost like a DNS tree but with address space. At this point I'm going to be really scared to show my face outside of this room at another routing thing because it's a very very stretched analogy. Basically, this allows people to have address space, to prove that they actually own that space. Then they can say, my AS number, my network, the unique identifier, my network, is this and I am tying these together. That way, if this was actually better deployed, kind of like DNSSEC, you both need to sign and have validators, if this was better deployed, people would have been able to look at the route and say, this route definitely belongs to Amazon, I know that this prefix belongs to Amazon, there's a bunch of

certificates that show it, but some other guy over there is announcing it, that's definitely not right. They would have been able to filter the route out and drop it. As I say, one of those emerging technologies that once it's better deployed would have saved us in this case. While we're on the topic of what you really should be doing, motherhood and apple pie, and brushing your teeth, there's a well known thing called BCP 38, which says you should only accept packets, much of what we've been talking about is only accepting routes, this is actually saying you should only accept packets from customers who should be announcing those. Sort of ties with routing information, if somebody is accepting a route and is allowed to, only that person should be allowed to be sending packets.

This is sort of a drum which we have been banging for many years, SSAC published a very similar paper which is SAC004 I think, if I remember correctly, which is largely saying if you're running an ISP and a customer starts sending you traffic, make sure that that customer should actually be sending you that traffic, here is automated means to apply it. Basically, if you're running DNS information, or DNS services, you should really be making sure that you understand how your routing infrastructure works, that you are monitoring your routing tables, otherwise you are going to show up, kind of like Amazon and [inaudible] shown up before. You should be verifying your

configuration to make sure that you're actually really still filtering like you think you are. Anybody who has run a network for a while, can spend a lot of time inserting filters, but somehow while you're poking at other things, somehow they seem to get lost, you turn them off while debugging something, you forget to turn them on again. Relevant SSAC publications, you'll take that?

UNKNOWN SPEAKER: Thank you Warren, here's a list of some of the publications, maybe even be comprehensive but as Warren already mentioned, SAC004 talked about [inaudible] and various other ones that we've put out over the years have touched on parts of these issues so that's a reference guide for you there, for some of our publications. Then I am going to wrap this up, of what you can do to help. Which is really just a review of some of the big things that [inaudible] and Warren just covered. A lot of this, awareness and socialization of best practices, which as we know is very difficult. We think that this particular incident brings some attention, shines some light on the problem and actually puts it into real dollar figures, as to the kind of impact this kind of attack can have, imagine if they'd set up fake websites for a lot more properties that were sitting on that same infrastructure at the time, the could have done a lot more damage. So, the BGP practices, MANRS, and prefix filtering. This is a great use case for

DNSSEC, and so that's... I know it's a subject near and dear to many people's hearts here, not just the signing it's the validation. That's what protects you from having these kinds of attacks, hit your particular customers, is the validation side. The validation, obviously, if you're customers, that's from an ISP perspective. From the sites perspective or the mail server that got hijacked or whatever it is that got their space... the cache poisoning attack from something like this would be... the perspective would be to sign the zone. Of course, you should be doing this anyways but not enough people do, as monitoring what the heck is going on with your DNS in the first place. Whether or not, you're running your own authoritative servers and are providing service, or you're actually going out to the web, you might want to run some of that as well, depending on... you're making sure you're going to the right places. Traffic monitoring, and the way your routing environment is actually set up. We've talked about different ways you can manage that, or have filtering in like... but understanding where their traffic is going to be coming from, and going to and should be. Mentioned the authoritative DNS servers, I'm in different places and this is... one of the ways you can actually make an impact is add this as part of your RFP, requests for proposals, around networking and things like that. What are their policies around filtering and all these other good things we've been talking about when you're making purchasing decisions, especially if you're doing things at

scale. Some of the folks here do things at scale, so that could be as a group though, powerful if you can actually get enough people doing that. Of course, monitoring both DNS and BDP for what is going on, and taking that to management, this is the soft and the hard one. This is why this, again, this case is so interesting because you can actually point to this and say look what happened. Now imagine that was our website and our customers were being redirected, what would the financial impact of that be? You have an actual case of this to prove it, it's not a theory, it's a problem with doing it in the lab or showing it at a conference, it's great but until somebody actually goes and robs somebody, management typically don't pay attention. Usually, they don't pay attention until they rob you, yourself, then they close the barn door afterwards. But anyways, those are some of the areas I think... the old saying goes, never waste a good crisis. Never waste a really cool security incident like this. So I believe, yeah we're onto Q&A, so... we'll take questions from the audience and if we don't get good... if we run out of questions, I see we already have one, so questions or comments from the audience or thoughts, and we'll talk amongst ourselves if we run out of time, and then we'll break early if we need to, we got... I believe we're scheduled for another 20 minutes, or 25 minutes, but we'll see how this goes. Go ahead.

UNKNOWN SPEAKER: Thank you. [inaudible] from [inaudible]. First of all, it is not Star Trek. Second, you mention all these things that people could do, if instead think of what people are doing right now, what's stopping the criminals using this trick to steal everybody's banking credentials, why wouldn't they do this instead of spamming the net with phishing messages?

UNKNOWN SPEAKER: So, in some cases they are already doing this. This is not the first time that this has happened, there have been a couple of well known incidents where large amounts of the internet have been routed through China or Russia, were some of the more recent ones, Iceland was another one. In many cases it seems as though those are potentially, more sort of large organizations doing this intentionally. I don't want to say nation state actors, because that gets into weird geo-political stuff, but nation state actors. So, some of that seems to be being done for more sort of analysis purposes, there have been a number of instances where specifically banks have had their BGP routes hijacked, and in some cases it seems as though it has been for large scale collection of this. Often banks don't really want to talk about that but there have been some relatively well documented cases. There also seem to have been a number of cases where people are just doing this in order to test that the capabilities work so that they can then leverage it in the future, but as to

what's stopping people doing this now? Very little. It is happening, it is just not being spoken about in public that much. Largely because when it does happen, people sort of pretend that it didn't, otherwise people lose faith in.

UNKNOWN SPEAKER: Also, if you look at the Twitter feed that you just mentioned, there are incidents all the time, it is a continuous feed, so it happens all the time.

UNKNOWN SPEAKER: Yeah, there's nothing stopping it from happening, right? And because it happens all the time, part of it is trying to determine, is it an error? Is it malicious? Or is it not? Is it just a configuration error? One of the things to also note that, this is sophisticated, right? I mean you have to understand how BGP works, you have to be able to know which route to inject, and then also the critical point here is that they were smart enough to understand that if you did the route hijack, if people were actually monitoring the IP address of the authoritative server, it didn't look like anything wrong was happening. It was specifically the cash poisoning where DNSSEC was one of the primary ways from a DNS perspective, you could have actually helped with this. When I think about the spam and sending, it is so easy to send bulk emails, from a criminal perspective, right, it's a lot

easier to do that and get a fraction of the people clicking on something that will then cause them to earn money, because it's all about money in the end. I think the fact is that the sophisticated attacks, are getting more prevalent and I think as one we're stating, malicious attacks that have used the BGP infrastructure have happened, nobody's talked about it, what this particular incident brings to light is that it may become more prevalent, so we have to start paying attention to how to mitigate.

UNKNOWN SPEAKER: Let me add a couple of points here, one is that it's just harder to get a hold of a router and inject the BGP. However, once somebody creates a service to do this, which will happen, you saw this cryptocurrency theft, it will happen and then it's going to become a lot more prevalent and we are seeing, there's a certain group of spammers that do very localized BGP hijacking, in order to inject spam flows into large email providers, and they do it very cleverly and very locally, but they clearly have the capability to turn that into some sort of a service or a kit. Basically, that's how they could do that, that's where you see most phishing and things like that, is all put into a kit and anybody can sell it and but it, and you know all that stuff.

UNKNOWN SPEAKER: I am getting more scared by the minute. Is this session being webcast?

UNKNOWN SPEAKER: A, I don't know, I think it is. But, I think that this is a well know, at least amongst [inaudible] it's a well known enough thing. I mean, actually, while we're talking I figured two things out, bogon is actually the elementally measure of bogosity, well done thanks. It was in December 13th, a bunch of well known financial sites, specifically, Mastercard, Visa, a dozen of other financial services were routed through a telecom in Russia. This is not the first one, that one got a fair bit of news, but there were other things exploding at that point, which I think is often a problem with this. There's always something blowing up in the internet and so sometimes these get drowned in the noise.

UNKNOWN SPEAKER: Thank you for the presentation, very interesting and really great use case for talking to management and so on, on both routing issues and DNS issues. My name is [inaudible], I am from [inaudible], Costa Rica, and we run both the registry and the IXP. So, it's interesting to have this case because it applies to both. Back in 2015, we enabled RPKI router [inaudible] validation at the IX, so basically, since the very beginning 2014, we enabled our PKI but the [inaudible] we enabled it in 2015, and that's

basically to discard any invalid announcements that we receive. Interestingly when we start working with the ISPs, that was the easy part, because they got the feeling that it was something important and we got everything signed. But, I do know... something interesting is that at this point the thing that has been the most difficult part is that critical infrastructure managers, let's say that way, or the companies that handle, for example, root servers and some of the biggest CDN's are the ones which are not signed with the [inaudible]. That is something that we have been working with and right now we have 100% of our ISP with all the raw sign, however, the not founds that we already have and that we have had since the beginning, it's basically on root servers that we host and also on the other side with the CDNs. So, it's part of a comment, it's something that raising the flag and says, creating infrastructure should be following these practices, and on the other hand it's a question on what do you think that could be done to motivate these administrators to go ahead and sign. Thank you.

UNKNOWN SPEAKER: I guess one of the things I wanted to mentioned is, I believe that Costa Rica is actually one of the first countries to be doing an IXP that does RPKI and drops invalids, I think that Ecuador was also towards the top. That was great, the fact that people were willing to step up and do that demonstrated to a lot of people

that this actually worked, and was doable. As for getting more stuff signed things, a number of the RIRs have made it a bunch easier to sign stuff, so for example, APNIC has a bunch of services where they will do a lot of the sort of tricky crypto stuff for you. As for getting it wider deployed, it is getting some deployment, not as fast as we would like, but things like MANRS, which is sort of an ISOC led thing, I think is helping because it's very strongly pushed towards it. I don't know whether it's actually a requirement, but it's definitely you should really really do this. As for getting it better deployed, Wes, you we're part of the BGP set design group weren't you? Kind of... there were a number of us, it's been going on for what, like since 2011, 2012. The actual design of the protocol has been going on a long time, it's only in the last year or two that many of the documents have been published, and things are still changing. Recently in the IATF it moved from a SIDR working group to now SIDR ops, which is basically the operational side of it. So it is being deployed, it is rolling out, not nearly as fast as we would like.

UNKNOWN SPEAKER: We are going to take a question from online.

UNKNOWN SPEAKER: This question is from Brett Carr. I think we need to ramp up education on this significantly somehow. We had an issue last

week where we saw one of our infrastructure prefixes hijacked. Communicating this to a tier one provider was surprisingly difficult, they didn't understand the problem even when pointed out to then.

UNKNOWN SPEAKER: I will take this. I find this somewhat humorous in a really sad way, because I know that workshops have been going on for at least 20 years to teach routing best practices and filtering, some of those slides that you see in there were drawings that I had in workshops that I gave 15 years ago. So, I am very concerned about how difficult it is to do this, as I was thinking through what Warren was saying, it starts from an enterprise also, knowing that OK, I am going to send to my upstream ISP these routes. The ISP then as it aggregates, that first hop should be the one that's responsible for doing the filtering. Everybody else also, but the one closest to the actual enterprise, right? So, I am starting to think why aren't even enterprises looking at, OK, you as an ISP can I trust you to make sure that you're routing my networks appropriately. I mean, it starts with everybody else also, but when I looked at BCP 38, part of the issue for me was, it's all about [inaudible] filtering, right? [inaudible] also from the enterprise should be happening. This doesn't just speak to this particular problem but overall we're not doing network hygiene, which for 20 years, we know our best practices. I mean you look

at SSAC004, which talks about the filtering, that's 2002, that's 16 years ago. So, I agree with you about the education, because I do think one of the things we as the community who have been around for a while and have talked about this for at least a decade sometimes two, I think we also forget that people change jobs and roles and there's constantly new people coming in. So this is new to a lot of people, especially folks that are on the operational level that might be the ones doing the configuration, right? So, I'm 100% agreeing with the commentator of the question, the person sending the question, that we need to just do continuous education on the best practices and what I've found is to show configuration examples. When Warren pointed to the [inaudible] site, one of the things that [inaudible] does, it actually has an automated feed, so it's not just the template, but you can actually subscribe to BGP feed that will help you automate the filtering.

UNKNOWN SPEAKER: I guess I will add something. So, this is somewhat of a question for the group. How many people here were really aware of this problem before? I guess, how many people... OK... and how many people have had to deal with BGP hijacks as time goes by? OK, that's an interesting set of statistics. So, something that we've been discussing in SSAC was having sort of a tutorial, kind of like there's the DNSSEC workshop type thing but also sort of

how routing works, because having your DNS work really nicely is important, having DNSSEC is important, but if the substrate that you're riding on, which is the routing, isn't secure and working right, it doesn't really help. There's turtles all the way down type problem. So, do folk think that more of this sort of information will be useful? How routing works, how BGP works, how you make sure that your announcements are correct? Is this interesting to people? No hands, I guess it's all boring... no-one. I didn't phrase the question well. Would people like to see more information on routing at things like tech day or DNSSEC workshop? Woohoo, you've got work to do. Related to an earlier thing, yes this needs to be better deployed, I had a quick look at the stats right now, currently of the around 770,000 routes on the internet, only 8.3% of them are actually signed, or have RPKI stuff which means that 90% of them aren't yet. For people who have been doing DNSSEC for a while, some of those numbers might sound depressingly kind of familiar. You have been waiting for a long time.

RUSSELL BEAN: Thanks, legs are getting tired. Russell Bean, I work for an [inaudible] ISP and also for ISOC. You don't think that maybe it's a little bit out of the scope of ICANN, I am not sure, do you think the RNO should be more involved? I know there's the [inaudible] that helps with this kind of thing, but it's kind of voluntary and

being an ISP, the finance guys came along and I used to be part of a [inaudible] and now we're not because we can save $500 a year not being members, kind of silly. Do you not think the RNRs could do a similar job, or should they be doing a similar job to the [inaudible], making it more compulsory?

UNKNOWN SPEAKER: That's a difficult question to answer while at a mic with camera's, yeah possibly. Yes, I mean some of them are, I mean like RIPE has a really nice routing database that people can subscribe to. I have some address space from a different RIR and I put all of my resources in RIPE because they make it easy and friendly and nice, which some of the other RIRs don't.

UNKNOWN SPEAKER: I think that's part of the issue, because some comply with it, some don't, so in the end it's not useless, it's useful, but it's not compulsory so you can't use it as a guide.

UNKNOWN SPEAKER: Yeah, you can't use the [inaudible] stuff reliably, and what makes it even more fun is some ISPs will hopefully register your routes for your, if you haven't done it yourself, which is this weird proxy registration and then you have absolutely no idea where the [inaudible] comes from or why. Which is a fair bit of

that is actually, or at least some of that is so the impetus behind things like the RPKI, you can't make people voluntarily register all of their routing information if people aren't forcing you to. So, having a system where you can actually validate the information, you know, through a cryptographic means, makes it something trustable and hopefully sort of rebooting the systems that people will now register. [inaudible], RIRs and [inaudible] and stuff like that is all a horrendous mess, it looks really hard to fix it. We'll just start again with something better, I think is some of the view, and I am going to so get beaten up when I walk out of this room aren't I? As for whether it's something within ICANN's remit, you know, could be ASO type people, or I think this is SSAC not really doing a... this is how we're going to make policy or anything, because we don't do that, this is more of a hey, this might be interesting to people, here's important background information if you care, go and poke those people now. OK, is that reasonable? Did that actually answer the question or did I get sidetracked?

UNKNOWN SPEAKER: Most of the solutions that are out there are like MANRS, are sort of best practice but not everyone does it, or they're in the private sector and that's obviously not compulsory and has a cost.

UNKNOWN SPEAKER: There's nothing that we can force on this, unfortunately, there's nothing that ICANN should force. This is more just a here a bunch of people who operate DNS stuff, let's give a quick tutorial on the side and see if we can get them interested.

UNKNOWN SPEAKER: Maybe if the right probes could do something as well, as we all seems to have those, might help.

UNKNOWN SPEAKER: The one comment I will make, which, you know, may not be so popular. We cause our own issues, and so if we as a community don't do quote, unquote the right thing we're going to see more regulation. So, that's one of the reasons why as we see more of these issues come up, we really as a community should work together to do the right thing because I don't think we want this to be regulated.

UNKNOWN SPEAKER: Alright.

JOHN: Hi, I am John [inaudible]. Could I complain about BCP 38 for a minute, OK. If you're single home BCP 38 is perfectly simple, but I've talked to a bunch of providers that say, they had the famous

problem that they have dual home customers, both upstreams allocate IP addresses and the addresses leak out through both interfaces, you know, and they through up their hands and say, we can't filter. This sounds like something that should be soluble.

UNKNOWN SPEAKER: So you have BCP 84 which I like to actually refer to as that deals with multi homing, and there is somebody... it's an IETRFC, those that aren't familiar, RFC are requests for comments, and there's a guy called [inaudible], and if Google for BCP 38 operational issues, there is an excellent document that he wrote about operational issues or best practices surrounding dual homed filtering. So, I prefer to use 84 because that was supposed to solve the multi homing problem, and you're absolutely right, because one of the things that people forget is that routes go down, you have backup routes. So, a lot of reasons why people don't filter is because they don't want to spend the time to think through their network to say OK, when everything is working these are the routes I am supposed to get and pass on. Now, if somebody's route goes down, they use a backup route. What is that? So, there is something called traffic engineering and I will emphasize engineering, because you do have to think about all the alternate paths that might be possible also, so that when you're creating your filters, that then

if somebody is starting to use a backup route because the primary went down, that you don't, by accident, shut that down.

UNKNOWN SPEAKER: In this case, they're not backup routes, they've got two live upstreams.

UNKNOWN SPEAKER: Actually, it is definitely doable. BCP 38 with multihoming is definitely doable. The BCP actually just says you should only allow people to send packets that they should be able send. It doesn't really go into the... and that you have allocated it. However ISPs are lazy.

UNKNOWN SPEAKER: The question is how do you know? My provider A, and my customer gets a random chunk of addresses from provider B.

UNKNOWN SPEAKER: So you can figure that out with IRRs and RPKI will tell you as well, eventually. The customer has to say, I've got this address space from that guy over there, and then he has to come along and ask you to please allow it. Usually what happens is that the customer doesn't think of that until they start trying to use the prefix through the other provider, at which time it doesn't work

and they throw their hands up in panic and shout at you. They just turn off all filtering because it's easier.

UNKNOWN SPEAKER: The comments I've heard are along the lines of, if you want route my packets I am sure somebody else wants my $100,000 a month to do so.

UNKNOWN SPEAKER: Yeah. Unfortunately, I don't know how we fix that. I am actually checking if there is BCP 35 filtering here, there we go.

UNKNOWN SPEAKER: [inaudible] my question is about the DNSSEC fragmentation issue. Because the DNSSEC will increase the package size and sometimes it will, the package will be fragmented. I am wondering if the fragmentation will be a problem for the servers to use BGP because the package maybe routed to different nodes.

UNKNOWN SPEAKER: Could you speak up a little? Sorry.

UNKNOWN SPEAKER: I'm asking if the fragmentation issue will be a problem for the servers to use BGP. Fragmentation...

UNKNOWN SPEAKER: Shouldn't be, depends on which part you're speaking of. So for BGP itself runs over TCP, it's a TCP based protocol so it takes care of...

UNKNOWN SPEAKER: For UDP DNS. DNS most runs on UDP and sometimes the...

UNKNOWN SPEAKER: Sorry, I thought you were talking about BGP SEC, sorry, I didn't hear properly.

UNKNOWN SPEAKER: If you are asking specifically about whether or not fragmentation will be an issue, because there is big packet sizes, I know that there's been a lot of discussion in operational communities with regard to IPV6, because there is a lot of operators that draw packets that have the extension headers, one of them being fragmentation. Now, I don't think that's a reason not to do DNSSEC, I think that's a reason to make sure as a community, we sit there and say, look we can't have these kinds of issues like this AWS 53 Etherium happen, we need to have DNSSEC, if

you're dropping the packets, why? What can we as a community do? Because we cannot say that we cannot implement the security best practices, that is the wrong way to go. So yes, there are issues, right? We need to see where are they. Where will fragmentation be an issue, with IPV4, IPV6. I know none with IPV4. I just know that from an operational level ISPs have been known to drop packets with extension headers, so we need to look at the ecosystem, because we want people to use DNSSEC, we want people to use RPKI, and if there's issues from an operational level, let's go and see how we can solve them.

UNKNOWN SPEAKER: So shall we limit the DNS packet size into a small size, maybe smaller than the NTU?

UNKNOWN SPEAKER: So, we're actually at the end of the session here, so can we answer that question offline. Great.

UNKNOWN SPEAKER: The only thing that I will add is that there is a lot of interconnecting items here, and there's a lot of folks that go to the IATF that come to the ICANN meetings, [inaudible] so, you know, I would encourage people to participate in the constituencies that have stakeholders in those communities.

UNKNOWN SPEAKER: Yeah, that's a good way to close it out, thank you all for your time and all the questions and we'll go back to [inaudible] for the next session.

UNKNOWN SPEAKER: Thank you very much. We actually were not scrapped for time, we lost one presentation so, if there is anything from the audience, I don't want to get, but I don't see this is the case so we can thank you very much again. Martin Boil and Donna Austin if she is here.

UNKNOWN SPEAKER: He or they will talk about the CSC review.

MARTIN BOIL: Good afternoon everybody. I'm Martin Boil and I am a member of the team that's carried out a review of the CSC charter, nominated by the ccNSO. The CSC was an organization, a committee setup as part of the IANA transition process, with the role of providing a direct input from the customers of the IANA naming function on and over the IANA functions operator PTI, and so for the ccNSO I see this committee as being really an important organization in ensuring that we as customers of the

PTI get what we need. So, if we can have the first slide please, have I? Oh, excellent. Thank you for that Kim. The purpose of the CSC, its mandate were deliberately set during the discussions on the transition to be very very narrow and bearing in mind what I just said about the role of the CSC, face-to-face with the customers of the IANA naming functions, we... this was actually quite deliberately done, and certainly when we carried out the review, it was pretty clear to us that everybody involved in the process saw this as being a very clear and important factor in the success of the process. Overall unanimity about that, this slide is entirely historical and looks at the process but perhaps the most important thing is the bit at the bottom, we have been out for a public comment period, we received 6 comments. Most of those comments were supportive of the major findings but there were some additional points that we have tried to take into account in the final version that is now been submitted to the ccNSO and GNSO councils. As I've just said, the narrow mission and scope of the responsibilities of the CSC, very strong support that there should be no change or expansion to those, and similarly there... membership of the CSC, very limited, very small membership, two people representing ccTLDs, and two representing gTLDs, as being the membership supported by liaisons from all over communities, but the important thing is to make sure that the customer requirements and needs are being addressed, and overall the liasons while they are not members,

have been actively involved In the discussions in the group and have up to now always been treated as equal in those discussions. The difference is that should something have to go to a vote, only the members, in other words those representatives of the customers are... have the right to vote. But that, i think we all see as being an unlikely circumstance.

There has been an issue on diversity, because with a small group like this, trying to get diversity is actually quite difficult, and the key factor for appointment to the CSC is about ensuring we have people with the right skills and knowledge and also the ability to be actively involved and this actually requires them to be on a conference call once a month with quite stringent limits on how many times they can miss that. However, and I think I cover this later, we have made a slight amendment because there is work going on in work stream 2 of the CCWG and their results will have an impact on this. Introduction of mechanisms to deal with change circumstances of appointed members, since the CSC was formed, two members of the committee no longer... ceased to be working for a registry operator, and that process was... that process went through to identify whether that person needed to resign from the committee or not, so we have clarified that in the report and we have had to make a minor adjustment on the duration, should a vacancy be fulfilled and I'm hesitating here because I think again that I'm going to be covering this later

on, yes I am. We've also looked at reducing the number of meetings per year, this is face-to-face meetings a year, and this is more a response to the change in the ICANN meeting structure and so we're saying at least twice a year, the thing being to make sure that those meetings take place when the relevant people are going to be available. There is nothing to stop there being three meetings a year, but we would like those meetings to be appropriate to the audience. There is... and this came as a bit of a surprise to us, that there was no defined meeting between the CSC and the PTI board, and was a general feeling on both sides that there was a lack of ability to have a strategic level oversight and view of what was happening within PTI and an ability to interact with that, so the review has come up recommendations on that. That there should be a formal expectation for that to take place. Service levels, and monthly reporting. Each month the CSC reviews the figures that come through from PTI and assures itself that service levels are being met, we expect that to continue, but that the monthly conference calls might in the future be relaxed, but we saw this as being an operational decision and believe that is a decision that is best made by the committee itself. Obviously, in discussion with the ccNSO and GNSO councils.

There were two, i think quite significant changes. The first one is about being able to make amendments to the service levels, at

the moment any change to service levels has to go through quite a complicated process, quite a long winded process. The intention is that the CSC should have the right to discuss service levels and to agree on making small changes to those service levels should that be appropriate. This is actually again a fairly limited overall scope because the definition is on small changes, but it would be on things like amendments where there is a new service being introduced and therefore some service level having to be put in to monitor that. Another example was where there was not enough data, originally to be able to identify sensible service levels and that this would allow experience to be built in so that the service levels were always appropriate to the function that was being done. In any circumstance, the change would appear with the ccNSO, or would be communicated to the ccNSO and the GNSO, and we would expect that there would be some discussion within those communities before going ahead with changes to the service levels, in particular, as to whether these really were minor and whether people were happy doing that. So, it appears important, it has the importance of trying to cut out red tape where red tape is not necessary. The CSC has fairly recently published its remedial action procedures, when the CSC was created there was an example of what would be expected in the charter, the CSC has negotiated with PTI a way of addressing issues as and when they arise. This is now been done and

agreed, and we now believe that this is an operational decision and that in future the CSC should have the right to introduce modifications to that, should they deem necessary. There is a specific case that should the PTI be replaced with a new operator, then there would be some expectation for a review of the remedial action procedures with the new operator, so there is wording in the charter that addresses that, and a general expectation that the RAP will be kept under review. But, this then becomes something for the CSC to work out without us having to define it in advance.

There were three areas of observation where this is not a charter issues. The first one on that list is the layering of reviews, a whole series of reviews taking place at around about the same time, with overlap, with resource demands on the community and a potential impact of the work of the CSC being interrupted by having to address reviews. So, we, in the main body of the report look at ways in which the ccNSO and the GNSO can try and identify as to how to address these reviews in the most effective way possible, which will then help them avoid unnecessary overlap and perhaps... I would hope a better use of everybody's time. The travel support in the initial call for membership of the CSC, the call made clear that there was no travel support available for CSC members. We have come to the conclusion that there probably is and in fact by not providing

support to CSC members, it will make it more difficult, certainly for some CSC members to attend those face-to-face meetings, to provide feedback, and discussion with the community. So, there is a recommendation that travel funding be found from ccNSO and/or GNSO as appropriate. The last of those points is one that came in as an input from the consultation and it came in from a ccTLD that identified that there were a number of organizations doing various parts of the work but there are a number of potential gaps between who is doing what. 3 minutes, I am fine for that... and the... while this isn't directly something for the charter, this is something that the report flags as being for the ccNSO and the registry stakeholder group to start thinking about how to move this forward so that we can be sure that somebody is picking up the various parts of those IANA related activities. So, as I said, the final report includes an update, amended charter, was published on the 19th June, has been submitted to the ccNSO and GNSO councils for adoption and hopefully ratification of this meeting and I have been... I have learnt that we have got a minor mistake in the text, so there will be a minor modification that we will put forward to the ccNSO and GNSO during the discussions over the next couple of days, so that we are ready for ratification of the charter. That is it from me, I'm quite happy to take questions from anybody.

UNKNOWN SPEAKER: OK. I will allow one question before we close the coffee break. Thank you very much for presenting a very dry topic very nicely. So, it's a 15 minute break and we'll convene at 5 o'clock. Jay.

UNKNOWN SPEAKER: Just briefly, I am not for Martin. So, I am the ccNSO representative on the customer standing committee that is stepping down later this year and I need a replacement. Hopefully somebody who understands SLA's in some depth and has some strong operational connection to a ccTLD. So, if you would like to talk about it, please come and see me. If not, please apply because we need a high quality replacement. Thank you very much.

UNKNOWN SPEAKER: I am asked to say that we also need one for the GNSO as well, RSRYSG but they have a different more complex process of choosing that person.

[END OF TRANSCRIPTION]