

RAW FILE

APNIC 62
ROUTING SECURITY SIG
BALLROOM 3

SEPTEMBER 9, 2026,
11:30 TO 13:00 P.M. IST

Services provided by:
Caption First, Inc.
P.O. Box 3066
Monument, CO 80132
www.captionfirst.com

This text, document, or file is based on live transcription. Communication Access Realtime Translation (CART), captioning, and/or live transcription are provided in order to facilitate communication accessibility and may not be a totally verbatim record of the proceedings. This text, document, or file is not to be distributed or used in any way that may violate copyright law.

>> In case you're wondering what the next session is, it's the Routing Security SIG and we'll be starting in about 90 seconds.
(Pause).

>> TERRY SWEETSER: Okay. Well, good morning, Mumbai. How are we all? Welcome to APNIC 62 and the Routing Security Special Interest Group forum. We have a really interesting program for you today. Taiji and I have been slaving over this program for about five minutes last month, but luckily we've gotten, you know, a good collection of presentations. And I'm just eyeballing, we've got the three in persons here, I can see Tom, Gaurav, and Anurag. You're going to hear from Tom about what APNIC has been doing with as spa and the RPKI updates in general from APNIC.

Gaurav's going to give us the local context for India, including talking about the ROV, which will be interesting.

Anurag is going to repeat his topic from the other day because we thought, you know, we want to have someone explaining ASPA. For those of you in the room who want to know what ASPA is, wants to know what AS provider authorization is, Anurag is going through that with you. That's a tutorial segment.

Then we have Marie and a Cheyenne remotely. They'll be talking about ASPA as well. How many ASPA talks is that? That's three or four. In case you haven't realized, the theme for today

is ASPA. And of course if you follow me on Facebook you would have seen today ASPA's the new black. Yes, a few laughs, okay, good. Thank you. I was worried about that joke. Sometimes it would land, sometimes it wouldn't.

Okay, all good. So what I'll do is I'll get our first speaker up, Tom. For those of you online, you will find that Tim is moderating on the Zoom. Please if we are running out of time, we will actually push Q&A to the end of the session. But fingers crossed we have that five minutes available to us.

But thank you, Tom.

>> TOM HARRISON: Okay, good morning. I've got three topics to talk about today. The first one is ROA coverage, ROV adoption, and ASPA adoption. Second one is ASPA deployment and third one is an update on changes we've made to the network health dashboard.

So we start off with ROA coverage for IPv4. Everything is heading in the right direction globally and regionally, which is great to see. But the standout figure here is in the APNIC region with 22% increase in the last six months.

For V6, things are similar but even starker. So again, each region is seeing an improvement or nearly each region is seeing an improvement. But again, APNIC is doing -- or APNIC region is doing the heavy lifting here and making the biggest contribution to that global increase.

If we drop down into the subregional figures, we can see that East Asia is where most of the increases come from or really nearly all of the increases come from for the v4, and the story for v6 is similar. We drop down further on economy basis, even though there are many economies that have seen improvements in the last six months, it's really China that has --
(Garbled audio).

The story of origin validation is more mixed, unfortunately. We have got an increase in the APNIC region and some increases for LACNIC and ROA and offsets in other regions that globally it's a wash. This is not a new thing. ROA adoption tends to progress better than ROV adoption. This is the story over time.

Moving back on to something more positive, ASPA adoption has been going pretty well. You can see on this graph that RIPE deployed the support at the end of last year. And then ARIN and then APNIC in June. Since that time is up has been increasing so great to see.

The subregional data in the APNIC region, since it's quite early days that probably wouldn't be so useful. But the [?] economy breakdown is more helpful. You can see here that India is at the front with 57 ASPAs already, so really great.

But even then the per economy distribution is pretty good. Like we've got lots of economies that are adopting these objects

early. And we think this is indicative of people's general confidence in the RPKI system. So good news.

Moving on to APNIC's ASPA deployment. We set that up in [?] with the first object created by real members on the 25th of June. And per those previous slides, adoption has been increasing steadily since that time. We have nearly 300 ASPAs under the APNIC [?] now, the vast majority under APNIC requested CAs, created by APNIC. But there are some others as well.

And that's contributing to all RIRs. There are two particular aspects I'd like to highlight. The first one is that we provide suggestions based on RIPE's routing information service. And you can emulate that behavior by using the URL on this page and dropping your ASN in as the resource query parameter. And then taking the set of neighbors and the ones that are tagged left, those are the ones we will suggest as providers in the interface.

And this correspond to what we do with route management where observations in BGP held to ground suggestions for creation. Because they're based on BGP data, it probabilistic and user have to review them to make sure they're correct. But they're more likely to be overinclusive than underinclusive. And including an additional provider that may be inaccurate is much less of a problem than not including a provider that's really needed.

So I think on balance providing these suggestions is more useful than not.

And the second aspect is the BGP validation. As throughout management again, when somebody submits changes to the ASPAs, we go off and validate in BGP that those changes aren't going to cause problems.

So if a route is going to move from unknown to invalid or valid to unknown, then we just show it to [?] suggesting that they might want to review what they're going to do.

These two aspects of the [?] have prompted discussion internally and among the RIRs. So we'd appreciate it if anybody has feedback on how these things work.

And the last thing I have to talk about today is ASPA support in DASH. That's the network health dashboard. It's available for all members at DASH.APNIC.net. There's a new ASPA status page which for your own ASNs shows whether an ASPA has been created for them. Also it shows the ASPA status for upstream, for instance, the upstream of your ASNs, as well as for those ASNs that don't -- that aren't part of your account but out of which your prefixes are originating.

And also the routing status dashboard pages that been updated with ASPA information as well.

So this is the ASPA status page. And you can see the ASPA details for each individual ASN and also the upstreams there.

And this is the routing status dashboard, and just at the bottom table you can see an additional ASPA column there which shows the ASPA status for the ASN.

And in the coming months, the information team will be working on extending DASH's ASPA support to include BGP validation and also hooking that into the existing alerts platform so that if there's a problem related to ASPAs, you can get a notification about that and go and fix it up.

And that's it. Thank you.

(Applause)

Questions?

>> TERRY SWEETSER: So if you have any questions for Tom, please approach the mic and name and affiliation, please.

>> [?] considering a case that I have [?] ASPA from my ASN and mention my provider, although this is not very [?] but in case my provider get changed and I have an ASPA object which is not including the provider, so whether I will get any alert either mail or something from APNIC or what?

>> TOM HARRISON: There won't be anything at the moment, but with the upcoming work to DASH, if you register for alerts there, then the appearance of that provider will cause routes to become invalid and you should get a notification about that.

>> Okay. Thank you.

>> [?] for RPKI. We had used Korea to add the -- an item of ASPA. Does that be sufficient for -- for the implement ASPA or we have to use APNIC specific API to implement ASPA? Thank you.

>> TOM HARRISON: It's totally fine to use krill or another delegated CA to add ASPAs. You don't have to create them via the API or APNIC. So that's fine.

>> TERRY SWEETSER: Thank you very much, Tom.

(Applause)

>> TERRY SWEETSER: So our next speaker was actually asking a question before me. Where has he gone? Gaurav. So you're up next. Thank you.

So Gaurav is giving us the update on the India situation, especially with ROV and ASPA. So thank you very much, Gaurav.

>> GAURAV KANSAL: Okay. So thank you. My name is Gaurav Kansal and today I'm on we half of NIXI I am sharing about the India's RPKIs. So with respect to the APNIC region, all 56 economies. My talk is mostly to what is the current status in India.

How I ask move the slides.

First things first, what NIXI. NIXI is not only the agency who is providing IP sources, they are services of Internet,

first is IX Internet, second is dot NCC, third one is [?] the body which locates IP resources to the [?]

I'm very happy and pleased to mention that in our Internet exchange we are running barred software. I think Maria's team is mentioning the software and she's giving that talk so thank you to Maria and her team for this beautiful work they're doing.

Yeah. Next slide.

So what we are able to see in this slide is like in last six years our both ROAs and IPv4 and IPv6 has grown from 0 to more than 80%. What this graph is showing is the growth of IPv6 ROAs was far more steep than the IPv4. And roughly two, three months into [?] our ROA growth from IPv6 from nearly 0 to more than 80%.

So this we are not sure what was the reason behind this [?] but it is probably related to one of the big Telco which has recently launch and implemented the IPv6 network. This probably may be a case because of this. At that time, our ROAs for IPv6 was far greater than the ROAs for IPv4.

Then we see a slight dip in IPv4 and IPv6 in ROA stuff. We're not sure this was [?] error or network things.

Next slide, yeah.

So we -- we measure -- if we talk about the ROA, ROA adoption or the implementation in India, in IPv4 we have 89% and IPv6 we have more than 98%. But this is -- these figures related to total IP space. Whereas, if we count based on the previous announcements, the IPv6 ROAs is much greater than that. From 90 goes on to 95. So what this means is in the first figures we are seeing that this figure is of ROAs implementation of the complete announcement.

Like a few big guys, a few big [?] 24s will have [?] in these numbers. Whereas, if we count each -- if we give a rate of one of each prefix, then the ROA and IPv6 currently will be 85% and IPv4 will be 85% also. Which is nearly [?]. So this with respect to ROA means that the big IPv6 is also implemented the ROAs with the smaller one have yet to implement.

Next is talk about ROV. So this is good figures with respect to our country. So we are nearly 1% ROV implementation in India. Whereas, the world's status is roughly 27%.

If we talk about RPK infrastructure, it is cascading [?] mechanisms. So if my upstream level ROV will implement the ROV, I will automatically get the production. So if I'm not multi-home, I may not even implement ROV if my upstream hasn't enabled the ROV.

So he with have five big Telco companies in India which are sitting nearly the upstream of every [?] in India. I don't know all the lines they have, but because they are nearly on the upstream of every smaller ISP in India. If these five Telcos

implement ROV in the network, we will get the advantage and our figures maybe is two more than 80% or 90% within a [?]

So this is not what we are special implement ROV, but this is low-hanging fruit which we can easily get our figures raised from 1% to 80% or 90%.

Okay, so one thing why I was collecting data for this presentation of mine, we see that our ROV numbers increase to more than 40% at one slide. All this collected from APNIC, so thanks to APNIC for providing this dataset and the public records.

What we figure out that ROV got increased to 40%. So we will appear -- we were very happy that maybe our big Telcos have done it. But I think it was issue in the collection of the data because then the graph came back to 2% after couple days. Because this is a sliding window and in a dataset of 70 days or 90 days.

Yeah, next is -- next question is why we have ROAs -- good number in ROAs and not in ROVs. So we tried and NIXI did away with all its members, they did interviews, and this is what we got from our efforts.

For creating an ROA or an ASPA, I just need to log in on my APNIC or my IRINN, click and my ROV will get created. I don't need to implement anything in my network for creating an ROA or ASPA.

So for ROV, I need to set up a validator and put that validator in my forwarding path, actual networking path. So ISPs are slightly afraid, like for ROA creation and ASPA, nothing needs to be done in network. For ROV they have to add something in the network [?]

Plus they know if there upstreamen abled in ROV, they will get production by default. Everyone is waiting for other ISPs to do.

Third thing which is sort of fear is because this is ultimately a call -- first call in my business should not get [?]. So every ISP, if they implement -- if we ask them to implement ROV, they think that consider a case where any service or [?] is staying behind which is ROV invalid due to misconfiguration. My client not able to reach that service. And the client of my competitor will be able to access this service.

So I will -- it will bring a bad name for me. So they don't want to take a risk.

Everyone wants others to do, and I should be last one to do it. But this fear, we figure out we did -- find out that how many routes within the [?] ROA are invalid. So roughly 300 routes out of 55 -- some 50,000 routes are ROA invalid, which is slightly .3% of total ROA -- total routes of India. So this fear is very less.

But these three things are what affiliates have raised to us.

So we all have done ROAs, but nearly all of ISPs implement ROAs. And what actually they achieve -- what I'm trying to show in my slide, considering I'm an AS 1, ASPA 1, I have implemented ROA. And the second ISP which is AS 2 has implemented ROV also along with ROA.

So by creating an ROA, what I had done, I had provided a safe routing infrastructure, safe Internet to my computer ISP, because if AS 2 is of my computer ISP and by my ROAs they are able to safeguard their network or safeguard their customers.

So ROAs helping computer -- with respect to ISP, not talking about the Internet this ROA or RPKI. So they're helping to save the network, not mine.

So if I want to provide safe Internet to my customers, I have to do ROV along with ROA. ROA will not only will not work.

Okay. So we talked preparing for this presentation I was pretty sure if I implement ROV I will be protecting my users. But then this hijack came into the picture. In the first week of September, they published a report wherein they mentioned a big service provider in Germany, Europe, was publishing an ROAs.

So in ROAs, what happened we can publish that if I was less [?] network, I can help route of [?] 24. It is in my end what I can do my ROAs in my account.

Then what we do, we clear the ROA with maximum length so that I can announce a 16, 17, 20, 24, whatever network -- more than 24 is not an BGP. So whatever route I announce, my route will be accepted. So this is what we see [?]. So [?] doing the same thing.

They were having ROA /24. And ROA was valid for 16/24. And now it's out. It gets [?] because more expensive propriety.

So within hours, maybe less than an hour, all that traffic, general traffic which would go to [?] was reaching to this attacker. So what attacker is running the first case?

Victim side. The victim side service provider but actual victim was some software update. So they get a valid certificate because what now this [?] able to do, they will try -- they will do ROV. ROV was there because the route was announced 24 with the same AS. So they get the certificate and then they publish the malicious payload. So this attack was -- this was the means for more than one or two days.

So implementing ROA and ROV will not totally safeguard you. You have to make sure that your ROAs should be matching with your route announcements. It should not be a loose end. Cloudflare has specifically done everything that routes ROAs which are matching with their route announcements.

There's many ISPs that the big guys haven't done yet.

Okay. Not covering this. Just one thing to hold. RPKI and DNSSEC both using cryptography to prove things. But the benefit with RPKI is even if I will not get an RPKI, [?] upstream, my networking will not fail because if -- if the -- if I not able to [?] my routes will move from relevant to unknown one. So RPKI will not fail. Whereas DNSSEC will.

As respect to ASPA, more than the 20% of APNIC is from India. Out of 49 of 60 are IRINN registrations. They are working hard to push that [?] to publish the ASPA projects.

Next is [?] which are published the ASPA objects. [?] NIC, and many others. But all big Telcos here.

So NIXI's implemented the delegated RPKI as a setup for delegated RPKI because we have to use -- because as we are under APNIC, we cannot in our system independently, we have to do it under APNIC. We have set up that system and we have tested the subdelegated to have a set of subdelegated service for affiliate.

So this test is nearly in the final stage. We have -- we will integrate this with [?] 2.0. One more important thing which we are [?] we have to pass from all five RIRs. Out of five RIRs, we have had specifically running on its local [?] thing. And what we have noticed, if we [?] after a few days, then the report fetching forget phase and we have to start again.

Because there's no [?] mechanism in RDP. If we do it after a few days, it gets failed. So our request to add into consider, because every validator has to face the ROV. For implementing ROV I have to have ROA. So at least with fetching [?]

Okay, two things. What we have seen. Few of the ASNs withdraw after implementing it. I'm not sure why.

Two ASNs as a provider, [?] has done this Internet exchange has to have a provider zero only. This is what the best case it is. We are missing the names of Indian IX with AS2 as a provider. One with AS2 as a provider. So APNIC answer this question.

What we are committed. We will implement ROV and ASPA in the IX in quart 12027. I'm told we're running a [?] the [?] doing the feature we just have to do our homework for this. We are trying to get two operators for ROV.

We also after the attack I just discussed, we are trying to talk to our members for going for quality of ROA means. They have to specifically have ROA which is matching with the route announcement, not next length. And we were planning to have 200 ASPAs by this meeting but because we were the host we weren't able to focus on it. Target is 300 ASPAs from NIXI side.

That's all from my side. Happy to have any questions.

(Applause)

>> TERRY SWEETSER: I think, Gaurav, you've done quite well speeding that up. If we have one question for Gaurav, anyone? You want to ask, of course. Introduce yourself as well.

>> Thanks. Continuing with the process of asking question just before your own talk.

So two quick points. You mention in the end of your -- around the end of your slides that operators argue they're doing ROV gives benefit to their other network, not their own network.

I have heard the same argument. You're right in putting that argument. But it's technically wrong argument. So I will always argue against it. Maybe you can go a little bit -- yeah, there's the slide.

So remember if you as a network is getting a hijacked route, is it your problem or is it the problem of people -- it's the problem of the network whose route has been hijacked? If you're learning a Google's route from a fake, malicious character, your own traffic would be a tiny 0.00001% of the Google's traffic. But for you going is 30%, 40% of the traffic.

So I would argue it impacts you more if larger networks are hijacked and you are accepting it.

>> GAURAV KANSAL: So this slide what I mention here is this slide representing to our members that you implemented ROA but not implement ROV. So unfortunately you are helping your customer. So you should implement -- you're helping your competitor, so you should implement ROV also along with ROAs. This is what our pitch to the members.

>> Okay, that makes sense.

And other point, you said the single home network do not need necessarily to do ROV. This is not completely correct because in India context there are many networks which are single home because of transit but they are [?] a lot. I know many networks, you can look around on the table you can find hundreds of networks which are upstream but they are IP exchanges. They would have hundreds of rofs. But not doing [?] still accepting a hijack from your peer side.

If you are going to exchange, you should still do ROV rather than relying on your upstream.

>> GAURAV KANSAL: Agree with you. I was just trying to make a point that at least my five big Telcos implemented, instead of having a one person figure in the chart, I would have 90%. This is the point is I agree that everyone should have to implement ROA. So this is what we are seeing. We have timeline of some implementing on AS. We will try to do it by --

>> I think the goal should be on network security.

>> GAURAV KANSAL: I completely agree. Thank you.

>> TERRY SWEETSER: Gaurav, thank you very much for your presentation.

(Applause)

And you stepped down prematurely. You need to go up there now and give your talk. One item I think came out of that conversation too was very, very important, for larger operators, if they do ROV, they provide something I call herd immunity because they are actually scrubbing the table for a lot of their AS code.

So if you're someone the size of Hurricane Electric where you're AS count is 60,000 ASNs and you're doing ROV, you're actually protecting 40% of the routing table doing that across your backbone.

So Gaurav has made a good point, top five operators in India start running ROV, you'll probably get to 80% ROV coverage straight away.

Thank you.

>> ANURAG BHATIA: Good morning, everyone. My name is Anurag Bhatia, I work at Hurricane Electric. You have heard them quite a few times. Original idea for this talk was to cover ISP deployment at Hurricane Electric, but spent significant time in explaining ASPA first and many people may not be aware of it.

Some of these slides will explain how ASPA works and the current state of route [?] and then I'll jump in on how we have deployed it.

So how filtering is done presently. At this point of time, it's -- it varies from operator to operator. Some base -- some generate static filters based on route objects. Some generate static filters based on [?] some based on email they have received and many don't filter at all.

Some of -- we do have some large back bones which have a policy like just filter until ten or 20 prefixes after someone is announcing more than 20 prefixes just open up the whole filter. So varies from operator to operator. It's very custom.

There's no standardized end-to-end filtering. And this is not just true in this context, it's true globally.

One problem with the whole filtering has been the RIR. These are Internet routing registries. The original idea of RIR was that you publish your intent in a public registry and then you do the same thing in BGP table. This would give an ability to go and match your intent with the table for the others and they can validate whether it's [?] out or not.

It wasn't around security. RIR carried this flaw where anyone can create a route object for your prefix across many of the open RIRs.

The other problem with the existing filtering mechanism has been tooling. There is no easy way to deploy these tools. Take for instance in the end I have BGP Q3, that's problem for for

doing the list for resolving AS number or asset. You get the prefix list but you need a mechanism to present this router. You need a mechanism to update this in your router.

Largely that part is missing and that be causes problem. It's safe to say that it's extremely broken at this point of time.

Last night I was looking into the address space on the -- across the Indian back bones. Almost 5% to 6% of routes which they're announcing they don't match the RIR intent at this point in time. It's there because it's hard to validate it at this point of time.

They cannot -- it's the same problem as ROV. They cannot block it now because it's running for years. So this has to be fixed. Which brings us to ASPA.

Why we need ASPA is the point. So ASPA stands for autonomous system provider authorization, and the idea for ASPA is it brings the RPKI tech into path filtering. RPKI ROV would only filter origin. RPKI ASPA would help validating the path as well.

Since it's coming from the RPKI, it carries the cryptography and the authoritative part. It doesn't have the same problem where anyone could create a route objective or asset.

In this case you can expect only the authorized holders via RIR mechanism to be able to hold that. It's way for secure than AS-SET. And because it's part of the tooling, it brings the same similar set of tooling where you can connect your route we are a validator and just drop the [?] coming in from there.

It is going to be -- once it is deployed at scale and once it is available at scale, it is going to be very similar tooling across various hardware vendors.

So here's an example of ASPA, it has something called customer AS number which is the number you are looking from. Provider AS number, it can be multiple if someone is multi-homed. It would show the RPKI trust anchor and which the network is visible. And one important distinction here is that it only shows you the upstream. You're supposed to publish only your upstream, into the your peers in the ASPA.

So on the right side you see a [?] I picked this up from the validators. On top you have AS 15538. AS 835, 6939, and 34927, so essentially they're signaling to the world that only they're allowed to show their route. No one else has the transit in this scenario.

So how the validation logic work. Validation is done across the ASPA. You can have essentially three states. Again, very similar to RPKI ROV. Whether it's a provider, not provider, or unknown state. You can call it a provider if you match the ASPA cord with the routing table. You may call it a record if it's not matching at all. Record says AS 1 is the upstream, or AS 2 says it's upstream it would say nonprovider.

And it would be for considerable time as they go along assigning -- creating that's records.

ASPA does not give agency whether it's upstream or downstream. Coming to that in detail with the very free [?] concept. But the broad idea there is you just declare your upstream and ASPA validation happens on each hop of the AS hop, the AS part, it does not care whether the next part is your upstream or downstream, it takes care of that part.

We have to see how this works out, but that's the broad idea. Right now Windows are still working on supporting it. You have something like reject not provider routes from your peers, from your downstream, and at some point from your upstream as well essentially to drop all invalids.

So here's a basic example. Let's say you got AS 1, it has got an AS 2 as upstream. AS 3 is the peer and then AS 2 has further upstreams. So you can typically expect something like AS 1 having a simple record, AS 2 is the provider and then AS 2 having record and AS 4 and 5 are provider. AS 2 will not mention 3 because it just appears in this example.

For AS 4, they would see the route coming from 2 and 1. If they're validating, they can easily validate this route.

They will know that AS 1 -- that this announcement is valid because AS 1 has marked AS 2 as it's upstream in the record and the validator will pass this info to the router and router will be able to match this.

AS 1 will see AS 3, 4, 5, so AS 1 sitting here will see 3, 2, and 1. And it will still be okay. It would still match it because by design, ASPA accepts one level of nonvalidation. This is important with the protocol. Now 3 has not matched 2 as upstream. This is clearly a demonstration where it says until one point it is still okay.

But if there is one more upstream beyond this point then it's a problem.

Okay. So route leak. Let's say AS 2 leaks a route to leaking of route of AS 3 towards its upstreams 3 and 4. If 3 has already published a record, the leak would be detected and 4 and 5 can easily block it, assuming they are doing the validation.

If AS 3 has not created a record, the leak will go through. So that's the -- that's the fundamental idea here for ASPA validation.

This brings us to the important concept of valley free routing. Because ASPA is taking care of very important distinction, which is that your peers can carry your route to their customers and their downstream customers, but your peers cannot announce your route to their upstream.

How this works out in real world, let's take an example. Let's say you're originating from AS 20 -- apps for using these

public tables. I had private ones and it was too long to say them. Last night I changed the slide to these shorter ones. So AS 20 and so on.

Let's say AS 20 is originating prefix. It goes to AS 21. 21 is the upstream. This is expected behavior in routing.

AS 21 is taking it to 22 which is downstream. Expected behavior. In Indian context, imagine it to be [?] any of the upstreams. Route is going to the upstream, coming to the other downstream. Okay so far. Let's say AS 22 goes rogue or goes crazy or as usual, does a mistake and leaks this out to other upstream.

This is a problem because 22 is sitting between 21 and 23. It's not the provide are for these. So this is a classic case of route leak. You're learning out from one upstream and leaking to the other upstream. It goes to 23 and further to 24. This is a bad announcement which will now go to 24 because 23 has got it and 23 will announce its stable too for its customer, 24.

So this right now the shape looks like a valley. This is something ASPA can detect quite well. The way it would work AS 20 would have record stating 21 as its upstream. 21 would have record stating something else as its upstream. If it does not have any upstream, let's say it's a transit-free, Tier 1 provider which good idea is you would put AS 0 as the upstream, then the -- then it would necessarily won't have AS 22 as upstream. AS 22 would have 21 as upstream, not the other way around. So anyone looking from far off location can detect that this is a classic case of leak. It's going up, down, and up and down again, a classic valley formation. And this would be detected and dropped. Assuming validation is happening.

Another case of valley free routing, now as the shape shows in this case, from 20, goes to 21. Let's imagine the relationship here downstream, upstream, and let's say these are just peers, no upstream or downstream relation.

And in this example there's no valley happening. Even though 21 would not have 23 as provider, still AS 24 would be okay in accepting the route, because until one hinge point it is okay. It is acceptable. Beyond that it's a problem. Imagine there is one more AS after this then it's a peer then it would drop.

You cannot have two AS ops in between. So that's the broad idea of valley free routing.

The idea is quite straightforward and simple and takes care of the modern architecture we are rough networks peer and have multiple upstreams, it should take care of that part.

Let's go further. I think I have limited time remaining. Go fast. How can you create ASPA object? Part of RPKI a few times, you do it in the same way. You go to your registry provider, right now you can do it on APNIC or ARIN and RIPE. APNIC and

ARIN support. APNIC launched 9th July, quite recent. Here's the APNIC portal showing the -- this is a pic from the support page. You can go and create a record there. Video record, you need to be careful to define all your upstreams. If you miss your upstreams, you will have problem. Maybe not right away, but you'll start getting problem. You'll have problem with us for sure.

Configuring ISP for route consolidation. You need the support in the hardware. It's supported in BIRD, if you're running exchange, please consider supporting it right now. If you're running Cisco or Nokia, they have stated they are supporting this year. It's a feature they're rapidly working on so you can expect support very soon.

How ASPA is configured. You connect the router with validator. You might have already done it if you're doing ROV. And then you learn the state. You enable the ASPA validation and then apply the policy for invalid state. You can roll it out in gradual scale. You can say I want to drop it in as peers as the starting point and eventually from downstream and some point even from upstream to be completely protected.

Route DLR inator support, this will give you regional validation as the [?] how we have deployed it. This is the original idea of the talk.

So in our case, it's custom setup. Hardware is hard to support most things. We have deployed large-scale filtering earlier, and we did it without our support of ROV. We did it in a system where we look at 25, 30 odd rooms and we go through that list and based on that we generate filters and we push them to our edge routers around the world.

We create a fairly large backbone, 350 plus cities, 55 plus countries. All these edge locations are sitting in these locations, they get custom filters pushed to them.

And when we detect a bad route, whether it's ROV mismatch or an ASPA mismatch now or even the older RIR mismatch, we detect that and push it to our edge routers.

So it's one the rules in this long list, I published a lot of them. The broad idea there is if it's ROV invalid, route mismatching we will reject it. If it's not an AS-SET we will reject it. This is limited to 60. Over 60 we [?] the route. We don't want 200, 300 [?] showing up in the routing table. It's usual good practices in the same system and ASPA has plugged into the same system and the routes which are invalid get pushed to the edge routers. This is the broad idea.

Also besides the backbone part, we maintain this public page, bgp.he.net and we have started covering the ASPA. This is from last night.

As of now, globally 3% of ASNs are covered with ASPA. 2900 out of 88,000 ASNs. The list is larger from right because right was earlier in supporting the ASPA. APNIC has recently launched so the number should go up are rapidly for APNIC as it has started.

Some key limitations. This is the last slide here. It has a network effect so it would be useful when more and more networks do it.

It does not differentiate between relation of peer and transit mix. And this is especially true for many large networks.

There are hundreds of networks around the world which would have a pairing relationship with other network in one continent while they may be the provider or downstream in the other continent. In ASPA you can only declare that other network is your upstream, not regionally. So it would be regardless of location and protocol as well. You can't have a specific rule for IP, before IPv6 you can't say here's my IPv4 upstream here's my IPv6 upstream. Whatever you have it defined, anyone would validate it.

Both ASPA objects and validation for it to happen it needs -- you need the validation as well on the other end. Which is the case with the -- which is the problem with ROV as well. Just by creating ASPA you will not get protected. People have to eventually start validating it as -- I give this popular example of check signing.

We wouldn't in situation where people were just throwing off checks to bankers and they were, let's say, just, you another, and cashing it without looking at submitter. Then we got to a situation where 80% plus of people are signing checks but no banker is checking it because ROA is not really happening. And the signing would appear pointless in if other side is not checking it.

You need to have other side as well, which is on the variation. And networks do need time because Windows need to support it. But ROVs, it's quite robust. So please consider ROV if you've not already done it so far.

That's about it. If you have any questions, feel free to ask.

>> TERRY SWEETSER: We'll take one question for Anurag. Hand up. Affiliation and name, please.

>> Hi, Anurag, this is [?] from [?] Communications. I don't have a question regarding your developing ASPA. How does it validate, or drop the package when it comes to some nonstandard peering protocols which ISPs have. Like they have a hybrid model of peering and upstream ratio. So how does ASPA affect -- does it drop packets immediately or what's the protocol there? The

peering and the upstream, suppose there's combination of peering with upstream with ASPA.

>> ANURAG BHATIA: That's the problem in the .2 you can defer to that. You can only declare a network as your upstream or not upstream. When you don't declare it you can't have this mixed relation. This is true for some large networks around. That's actually a problem to some of the large networks.

>> Is there a path forward which you think that can be modified and, you know, made better or some protocol that added to it or how do you -- how do you --

>> ANURAG BHATIA: I think the -- a lot of these things are discussed quite a lot in various other forums. But yeah, I'm not really sure how one can fix this sort of thing.

An important distinction here, even if it was supported, I doubt many people who have this complex relation would want to document it in public. People know it when you look carefully at the routing table or when you learn it word of mouth. But I don't they would publish that, especially large big networks.

>> Thank you.

>> TERRY SWEETSER: Thank you, Anurag.

(Applause)

>> Quick question only. So you told that [?] checking for ASPAs, not only publishing. Have you any list of who ASNs around the world have implemented the ASPA checks?

>> ANURAG BHATIA: I don't think that's there. Probably one has to measure it in the similar manner as the ROV measurement APNIC is doing. Maybe good idea for APNIC to add that to check as well.

I'm aware of Cloudflare doing a bit where they intentionally put ASPA mismatch and then they look around various collectors on who is getting it and who is not getting it. So they have that check enabled.

Thank you.

>> TERRY SWEETSER: Thank you, Anurag.

Our next two presenters are remote. First up will be Maria from NIC.CZ. I assume they're up and ready to roll. Okay, five seconds.

So those of you who are not familiar with NIC.CZ, the people who actually publish the BIRD software and I think if there's a parent point or UX that's not using that, it's going to be a very rare thing indeed. So when Maria is ready, we'll get -- technical difficulties.

Ah-ha, there we go. Sound check.

>> MARIA MATEJKA: Hello.

>> TERRY SWEETSER: Excellent, thank you. Go ahead, Maria.

>> MARIA MATEJKA: Okay. Hello, everybody. This is Maria here and sorry I wasn't able to come in person this time because of [?] capabilities.

I have a problem with BIRD Internet routing NIC.CZ for more than ten years and I'm going to finalize the ASPA series with my favorite kind of topic, which is what can possibly go wrong.

I've been a developer my whole professional life. I've never been a [?] but that doesn't stop me from trying to break things. And when it's a technology that's expected, it made the Internet routing more pleasurable, it's the ideal candidate for trying to find problems.

And don't get me wrong here, please. ASPA's a good thing and you shall deploy it. But to help you --

>> TERRY SWEETSER: Can I just interrupt you? Sorry. We've got a malfunction with your slide deck. With you be able to share your slides remotely?

>> MARIA MATEJKA: Like this? Is it working now?

>> TERRY SWEETSER: I think we've got to change the view here and then we'll have your slide deck.

>> MARIA MATEJKA: Okay.

>> TERRY SWEETSER: Apologies for the interruption.

>> MARIA MATEJKA: I hope it's working through the Zoom.

>> TERRY SWEETSER: Zoom is giving us some issues. It's still sitting there saying started screen sharing. There we go. Perfect.

>> MARIA MATEJKA: So now it works? And now it works.

>> TERRY SWEETSER: It is working beautifully, thank you.

>> MARIA MATEJKA: Okay. So I've been a developer my whole professional life. But, well, I sometimes QA. If you wanted to deploy ASPA, well, this is the talk about the [?] cases which are not for the specification or whatever else. And there are some mitigation options so that you don't get called by the possible progress.

And it should be set up wrong. If everybody configures everything properly and everything worked properly, almost, because of hardware failures and so on, it's the ideal start of degradation. It's actually a situation where somebody screws up.

And it's about how to get out of that. The most obvious situations are when it's myself who made a boo-boo, that's, well, often you have to deal with problems which are caused by somebody else. Or worse, somebody has to determine which configuration is to fix at which operator to rectify the routing. And it may not even be a connection between you and somebody else. It may be a [?] part of the work.

So for now, as we already heard, we have some places to configure. And to make things more complicated, with ASPA, we are adding more places where one has to configure something.

There's obviously ASPA record in the RPKI. But you should configure BGP which is point us to the [?] not the RFC, because ASPA only validates what you received and doesn't validate your immediate neighbor.

And BGP rules also add the customer [?]. We should not forget about all the local policies and also as was already said, the IRR AS-SETs. And with all these things in one pile and I think I forgotten something, I highly doubt that if I ask anybody to configure my routing that they would do it right on the first shot.

Even if I tell them make no mistakes.

So what to do. The aviation industry has a long history of preventing disasters. And if you look under the hood, you will see a system which is extremely resilient against screwing up. By having multiple safeguards and layers of redundancy.

And the whole principle is often called a Swiss cheese. You have holes in every single slice of that cheese. But if you layer enough of them, nothing gets through, which also means that nothing ever prevents the disaster completely. But one must often screw up a lot to cause an accident.

So after this long introduction, we can dive into the ASPA deployment. We are doing this for educational purposes to avoid or recognize such problems in production.

As was already said, signing is much easier than verifying. Because well, networks started to sign the ASPAs, but the verification itself needs some router support. And deploying things to the router is quite hard.

So this brings some additional specific situations, like when ASPA becomes invalid [?] it gets propagated far, far away without anybody stopping that route. And that would be technically okay at sometime when it gets finally dropped.

Well, it's leaked, it should be dropped. Well, it isn't. And I have no data how much ASPA invites are actually harmless and how many of them are leaks. And that doesn't matter. In the end it doesn't matter, at least for now.

Because what matters is when you are signing the ASPA, you may forget, you may forget to include who are your providers and you have already here a question in the network is complicate. Well, the draft to BRIC does say that you must include them. No single must in an IRC has ever stopped anybody from accidentally violating it.

So you want to check all your BGP neighbors. They should be out anywhere in the world to propagate your routes for any customer code and the customer is the provider. Which means if they're your provider, you're the customer. This basically matches the definition of the provider and customer from the BGP rules RFC. But in the end, yes, anywhere in the world that ASN

is allowed to send your routes to the whole world, they are your provider and you shall include them.

But let's get back to the BGP rules. And BGP rules are actually the first thing which can catch your configuration error. You have to check that your BGP role is set to be customer, provider, whatever, which checks the other side if they have set their BGP role as well.

And you can check whether your role is kind of matching or at least not inconsistent with your ASPA. This is the first layer of defense. This is not complete. This is obviously Swiss cheese.

And your provider may be doing very weird things with the ASPAs, so the [?] number may not be the same as what ends up in the ASPA later on. But you can also check at least your ASPA verification after you prepare your own ASN. Because the ASPA verification runs on the route you have received. Not on the route after you send it.

So if you do this additional verification, you get some more checks. You will see whether the route still is valid after passing your [?]. And if it suddenly becomes invalid, there is something very seriously wrong with your setup. This may run even outside your border routers or router [?] on the route collector which actually working [?] and you can bring some red lights in the monitoring and force the -- force the administrators for network operators to do something about that.

And in the end, you can simulate the whole outcome of dating your provider set. Validating your own router flexor table with [?] these two approaches by themselves can save a lot of your time and help before your routes ever leave your network.

And if you have just little knowledge of process levers, you can check even more. You can look at what's going to be [?] by your lever. You can check the resulting ASPA in the way it's going to be propagated by that -- well this looks like a stretch. That's considered that there are some [?] sessions. Symmetrically you own a [?] route and you receive no route from that.

In that case, this layer of the Swiss cheese may save your day. This is an interlude, there may be some Swiss cheese in the picture because this is a train in Switzerland that the train is not going to helping you today.

We are looking at catching the error remotely. Let's say you can deploy any of those or it failed for some reason because complicated networks. We're getting in the risk of propagation.

You can check [?] but there may be routes which never appear down and ASPA verification is going to silently fail them maybe for months. It's going to be -- it's nice to see that NIXI is actually going to implement ASPA verification in their exchange

point. It's extremely simple with BIRD, you just write very simple code, [?] check, and you check ASPA upstream and I would like to please the desired outcome is dropping in [?]. And the reason is very simple. The route which is received from an exchange point gets often some preference over some transit.

And if somebody leaks routes into exchange points, that may cause large outages. And, well, at least now when the ASPA verification deployments [?] stopping this.

We would also like to get the OTC-invalid routes get dropped in the ISPs. But they need it to be displayed which is not supported. We have a fixed process and we hope to manage it successfully into the next minor version. But in the end, none of this is actually a good Swiss cheese slice.

All remote checks for a [?] to actually propagate the route to find out that it's bad. And this brings me to the program of the forgotten backup provider.

Because well, maybe there is an obscuring which works only in a case all your regular transits are down. And there is probably just been or two people ever knowing about it in your company.

It happens all the time. If -- and if you're asked the provider is created by somebody else, this backup provider is not going to be included.

Or somebody makes [?] number that also happens all the time, you can see it in the -- in the router configurations.

And with the backup provider, you won't notice the program in any public service. The routes simply don't show up anywhere. Nobody can notify you. Even if you're a provider, they won't catch it unless they do additional checks as I have on slide 8. Only when this specific backup provider gets activated, at that point of time your routes suddenly become invalid because they're propagated.

At that time when you expect that your network will be reachable at least through this backup, nope, it's going to be unreachable from [?] parts of the world. That's basically the major part of my talk today.

Then there is the question of the validation of the verification, because invalids are not equal. Let's get back to that. With your customers, it's very easy.

They basically should have the ASPA in order and you have contact with them, you can call them if they send you invites.

And they are your customers. You have a contract. You can reject that. They should not send you garbage. So getting things from a peer is a little trickier. At least if you have direct [?]. You probably still have some people who are in the link and they can communicate also about the ASPA directly and fix the programs.

And in an exchange point, I'm very much recommending to drop the invites on the route server because also that's where the public looking glass is and you'll see the program.

And then there is transit. And in a regular scenario, I expect that these routes have the lowest local preference. They get considered last anyway.

But if all of these are all invites, we are getting back to the question of the harmless leak again. Because the users often expect things to just work.

And the provider is to blame. As I said, you most definitely have some reasonable technical contact with your customer. There may be somebody to leverage in the future so that they have their ASPA right. And you can force them, that they force it on their customers and so on.

If the route is transit, it's way harder. You have to first - which very part of the ASPA is actually the broken one. Which may be in multiple points. You may need to [?] companies. They may be different time zones, they speak different languages, they may consider ASPA invalid program for another network. Of course this only works if you find some up to date contact information at all.

Don't forget, amid all of this mess, your sales department sends you a message every hour requesting you to fix this exact connectivity for a major customer who threatened to leave.

But in the leak is [?] I don't know. So this is my temporary recommendation. The ASPA [?] is very indifferent when saying that you should drop invites. But as we have seen there are some periods and that's why I say it may be better to depreferance on Val lids from your transits and use them as the last resort. This will become less relevant with more networks deploying ASPA validation. And as soon as your transits do it, you can start dropping invites as well.

But both parts stay. And make sure you don't send garbage out and whenever something looks wrong, look at it.

The no garbage rule applies when changing things in networks. If you're splitting marriaging networks, ensure that your routes stay valid after remembering basically by adding providers. It's perfectly valid for multiple ASNs to see each other as their providers. That works, that's okay.

And if you have five ASNs and routes flowing between them, they're providers to each other, it's okay. And you shall have automation. If you're ASPA is managed completely manually, you may [?] over time as RIR. Having one source of truth instead of automated checking or the constraints is greatly reducing the overall risk.

And you should remember that things change over time. All the checks should be reactive and regular. Otherwise, they become

outdated and useless. You should not check things once. And this includes checking for each ASN in your provider list whether it is still an actual provider. Because it's in the end a public authorization for them to announce all your routes.

And this brings to the end of my talk. This slide discloses how much LLMs were involved in this talk. And here are various contacts so that you can drop me a message or follow me on social networks.

And that's all. Thank you for being here and will being and now it's probably time for questions.

>> TERRY SWEETSER: Excellent. Thank you, Maria. I think we'll give the audience one chance for questions. We don't have any takers. Well, contact details were up just before, the decks will be on the website if you want to contact Maria or ask questions about ASPA tooling.

Thank you, Maria.

We'll move on to our next speaker. So this should be interesting to see how Zoom performs now. We're good? We're waiting. A hand up normally means five seconds, not five minutes. Five seconds. Okay. One second. Just negotiating. They're saying one minute, I'm saying one second. We're having slide issues again. You have slides?

So if you have slides, we'll drive and Shyam can talk.

>> SHYAM: Hello, can you hear me?

>> TERRY SWEETSER: Yes, we can. We're waiting for video to come up on screen so we can see you as well.

>> SHYAM: Can you see my screen?

>> TERRY SWEETSER: We can. Okay, we are good to go. Let's do it.

>> SHYAM: You can see my screen, right? Hello.

>> TERRY SWEETSER: Yes, we can. We've got it running very nicely now. Thank you.

>> SHYAM: Hello, everyone, I'm Shyam and I mostly work with Internet measurements on Internet security and [?] service mitigation.

Today I'm going to present our work on exposed BGP routers. In this work reasked simple question, how many routers with [?] protocol, which stands for -- protocol are publicly exposed to the Internet.

We may all know that BGP runs on TCP179, recognizing the importance of BGP and RFC 7454 recommends to put TCP179 to configured peers only.

However, there are situations, there are many routers who accept [?] followed by open messages from any push from the Internet. Then the routers respond to those unsolicited requests with either open masses or notification masses.

So in the case of open masses, it contains many useful information such as ASN number of the router, BGP configure the router and many other personal parameters.

While in the case of notification masses, the routers raise an inner message and tell the host to terminate the connection. So you know definition, we define an exposed BGP router as one which accepts the connection and open messages from any push from the Internet.

Take exposed routers we developed a tool and note it as scan-179. As you can see on this slide, this scan-179 consists of four main stage. The first one is scanning the entire IP address using Zmap.

[?] is to find potential routers from the Internet. Then we name those responders as TCP/179 responders. After that we created our own custom BGP speaker. That speaker contains all the parameters that are required for -- that are compliant with BGP.

For example, AS number of [?] like [?] and many other things. So we said that BGP open masses to TCP179 responders.

Next we collected the responses from those and we analyze the -- we collect the full packet and analyze the packets to infer routers, AS numbers, and the characteristics of those routers.

Following permission, we did one-shot Internet-wide IPv4 measurement in February of 2026.

Let me give you a brief summary of the responses that we got from our scanner. So we got around 6.25 million IP addresses that responded with Zmap 179 scan.

Then we got three categories of responses. In the left most hand you see the BGP responder which is 5%. So BGP responders are those hosts whose either respond to open messages or notification only messages.

We receive higher percentage of open responders, which is 51%, than notification only responders.

And the next two categories are known TCP responders and pre-TCP responders. NonBGP responders are those host who is allowed -- but after that they save messages which are [?] or didn't respond at all.

And pre-TCP host which means the Internet allow the TCP connection at all.

So ideally for secure BGP router, we expect the routers will fall into one of those two categories. Non-BGP and pre-TCP categories.

In our case, we furthermore analyzed open responders for [?]. Then the next step is identifying who are the ASes managing those routers. For that, we look into the open messages of those responders which say ASN that is configured is open or whether the ASN is public or not.

If it is public, then it is [?] to map those IPs to the ASN. If it is not, if it is nonpublic, which means either Azure, [?] or or private ASN, they rely on IP code to find the associated ASN.

In the case of notification responders, it doesn't contain AS number or any notification. In that case we Gwen relied on IP input to find the associated ASN. And we [?] those ASNs as associated ASN because we can't state those are the associated ASNs.

The next step is grouping those IP addresses to routers. Because generally [?] has a large number of interfaces. So when we get responses from those IPs, they might belong to the same router.

For that, we use the combination of BGP IG and ASN number in the open messages. It is because the CPIDs unique for router in ASN.

For the case of public ASN, it's easy. That key works while in the case of non-public open ASN the same -- the same ASN, which is private or [?] will be used by multiple other ASNs and the Internet. That's why we use the combination of the BGP, open ASN and interface ASN to group those IPs through to a router.

We furthermore characterize those routers into their operational roles as internal and border.

So internal routers are those routers who participate in the global Internet. It means an ASN might have -- might have their own ASNs for private use which is also known as AS configurations.

So we say we do the ASN is nonpublic. If it is nonpublic, then we classify those routers as Internet routers.

Then we have another Class of Router which is border router. A border router is one who participates in eBGP peering. It means -- it means it is having a BGP [?] with another ASN. The first [?] is that it should have a public ASN and one of its interface IP must belong to another ASN.

The [?] is best [?] personal practice where the two point-to-point connected routers have a sublet that comes from one of the participating ASNs. However, we found another category, actually this is the case where we couldn't confidently classify the router as border and internal and named it as unclassified.

In this category, routers have public AS number, open messages. But all interface IPs map to the same ASN. So it could be a router having an iBGP peering or it could be the host -- it could be the case that the host AS is providing the IP to its connected peers.

So using those categories, I present summary of our findings.

As you can see here, there were 141,000 exposed IP addresses which correspond to 20,432 unique routers. Among them, we found

large number of internal routers, then the border routers. We found around 16,100 internal routers. And those internal border and unclassified routers are managed by 4,194 ASes on the Internet.

Then we took deep dive into the APNIC to see which are the countries having -- having those ASes.

So as you can see here, Indonesia has around 387 ASes, having a large number of ASes, having expose the BGP routers. Followed by [?]. And Taiwan stands on the top -- on the list of the top ten countries in terms of ASes routers.

The next thing we measured the crit cattle of those exposed routers. For that, we have two categories. First one is border routers. To measure the criticality of border routers, we used [?] the number of connected ASes. Indicate the critical IP. So you can see in the figure here that 70% of border routers, 22 one AS, whereas, we found a border router that connects to ten other external ASes. It means in that router because of its exposure can affect ten other connected ASes.

Similarly, we measured the criticality of non-boarder router. For this case, we used the number of IPs configured in that router as an -- as an indicator of the criticality.

It is because number of IPs correspond to number of softwares in the router. Higher the number of IPs, the higher number of softwares. So that can [?] subnets.

We found that 55% of those routers have only one IP. And about 85% of those routers have fewer than ten IP addresses configured in those routers.

There was one border -- one non-border router which had 1,208 IP addresses. So we can imagine that this router was a particular router in AS.

To summarize our findings, we found a non-trivial number of BGP routers that are exposed violating RFC 7454. And those routers -- and the APNIC context 4,611 routers managed by 1,343 ASes. Among them 347 are border routers. But there are many interesting resources in the paper, for example, the vendors of those routers, diversity with Internet and AS.

And we talked to a few operators who know about the exist ambiguous of open ports. After 15 operators we asked, only two of them responded. And one of the operator in the [?] they were happy to get this information and they realized the problem in their router and fixed it.

While there wasn't another operator who said that this protection is not necessary because by default BGP's RFC 471 sends [?] to [?]

So with the audience here, I would like to know more about the existence of trust exposure and also that you can take the

routers and if you need any help with checking whether you have exposed routers, I can help with you also.

Thank you very much for listening to my presentation.

(Applause)

>> TERRY SWEETSER: Thank you very much, Shyam. That's topical because I'm aware that a lot of the [?] are critics. So those of you going home to Bangladesh soon or back to Indonesia, please check your networks because that's a very interesting routing security problem, especially if someone decides to syn flood your router because they don't like you.

Unfortunately, that brings us to an end so we can't do any more Q&A. I'll do the stand-up announcements. We have a list on orbit, please subscribe and participate.

Also the SIGs have a joint session coming up today where we're talking about SIG guidelines. If you want to have a matter with us about the policy in that particular session, please come along to that later on today.

And of course given that we're all tech muffins, IPv6 is after this, so definitely go in to see how things are going with that as well.

And if you're not aware, we are currently having co-chair elections on, so if you have not seen your email yet or you've not voted yet, please engage with that process because we would like you to, you know, engage with the democratic things that happen here at APNIC because myself and Taiji and others were all elected by you. I'm trying not to campaign on behalf of anyone right now, but it would be good if you went and voted.

Thank you.

Okay. I believe we can now close off the meeting with, of course, my great gratitude to Taiji who has been basically the workhorse of organizing this -- organization of this SIG this time around. And of course the APNIC secretariat who always do the heavy lifting, because we ask them too, but the fact that they provide all the AV and support. And we've had the AV guys here as well. You can see they're running around with the cameras and doing a great job.

To that, I will now close the session. Thank you very much.

(Applause)

(Session end at 12:30 p.m. IST)

This text, document, or file is based on live transcription. Communication Access Realtime Translation (CART), captioning, and/or live transcription are provided in order to facilitate communication accessibility and may not be a totally verbatim record of the proceedings. This text, document, or file is not

to be distributed or used in any way that may violate copyright law.