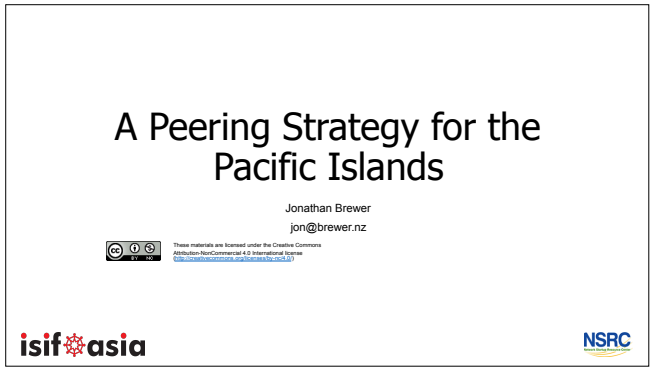

Kia ora koutou. I'm Jon Brewer, a network engineer based in Wellington New Zealand. This talk discusses preliminary findings from a larger body of work in progress around connectivity and Interconnectivity in the Pacific. I'd like to thank the University of Oregon's Network Startup Resource Center for funding some of the early work discussed here today, and the Information Society Innovation Fund for work since June 2015.

I've been based in New Zealand - a hub for the Pacific and Pacifica people - since 2003. In 2013 I hooked up with the University of Oregon's Network Startup Resource Center help teach network engineering in the Pacific Islands.

PacNOG is a group of network operators organised by the Pacific Islands Telecommunications Association. It's main function is to serve as a biannual gathering of network operators for workshops on a range of network design, management, operations and security topics.

These meetings are an excellent way to meet the engineers running networks in the Pacific, learn how their networks are put together, and understand what they need in terms of resources so that they can provide a good service to their customers.

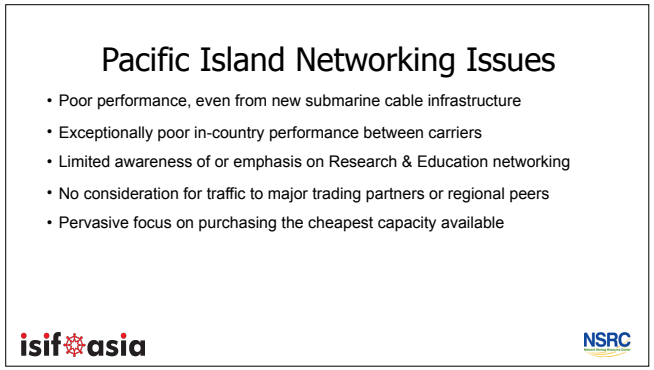

Through the six PacNOGs I've attended I've found a range of problems appearing again and again, in every country or territory we visit.

Some of these problems include:

And so, a year ago I began a project that I hope will help discover some of the reasons operators are having trouble, and provide hard data for engineers, managers, and regulators to consider when making decisions relating to Pacific connectivity.

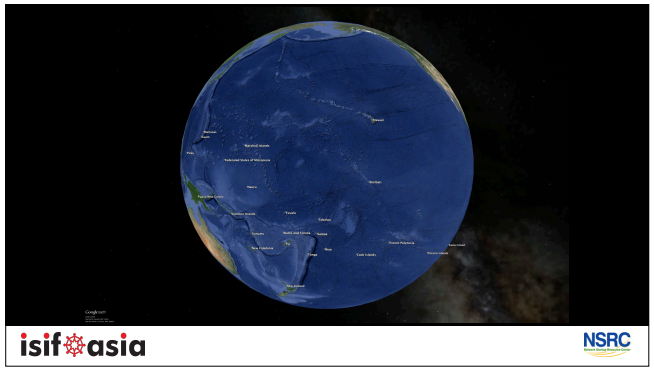

Let’s look at how it’s all strung together first. This Movie shows how connectivity has evolved over the past twenty-eight years, beginning with TPC3, the Pacific’s first fibre optic cable, and finishing up with 03b’s ramp-up in 2014. You’ll see years in the upper left-hand corner, and cable names in the upper-right. When you see a cable name in dark grey type, it means it was removed from service.

This movie condenses down a twenty-minute long section of a talk I gave around the Pacific last year, in New Zealand, Japan, Samoa, Australia, and Guam. Want the long version, buy me a beer & I’ll be happy to go into detail.

Taking the larger view - with cables and satellites in place - the Pacific is extraordinary well connected. Its major hubs Hawaii, Fiji, and especially Guam, have the potential for fast, low latency connections to the entire Asia Pacific region - and North America - and soon to Europe via a polar cable network.

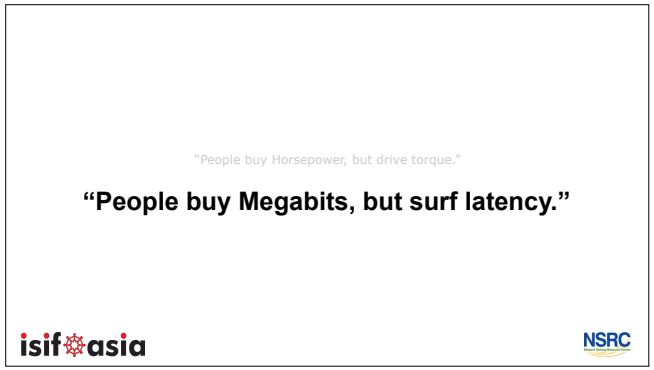

There's an old saying about how people buy cars. Headline figures of horsepower are great, until you put your foot down & spin your wheels, not going anywhere fast.

I like to think it's the same thing with megabits and latency. You can have the fattest pipe in the world, but if you're at the end of a high-latency connection it's going to be frustrating to use.

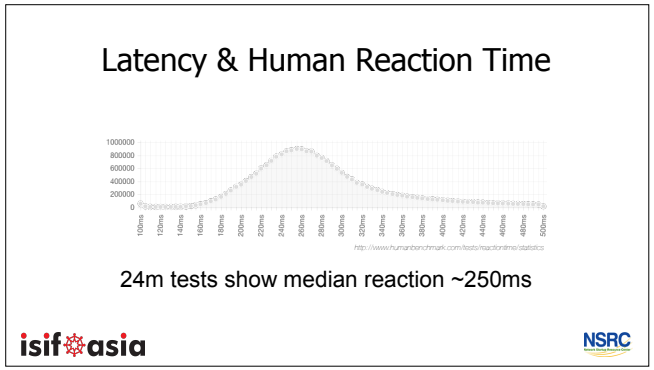

Most people can perceive a stimuli and react in less than a quarter of a second. When there's more than a quarter second of delay in a real-time application like Skype, Facetime, or distance learning, people notice. And it's not just real-time communications that are affected. Web pages today are complex packages of data, often taking a couple seconds to load even on low latency connections. High latency equals frustrated users.

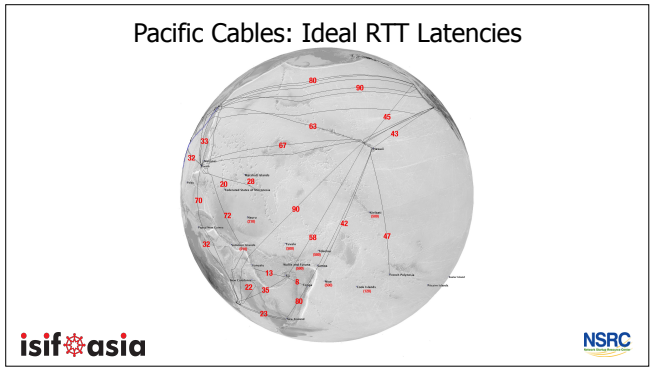

So let's go back to all that new cable infrastructure. Talking to cable operators & measuring cable segment lengths I've constructed a simple model of network latencies to look at how all these cables should be performing. Numbers are RTT in milliseconds. It all looks great. It's hard to find any combination of paths that would reach the quarter-second threshold for human perception.

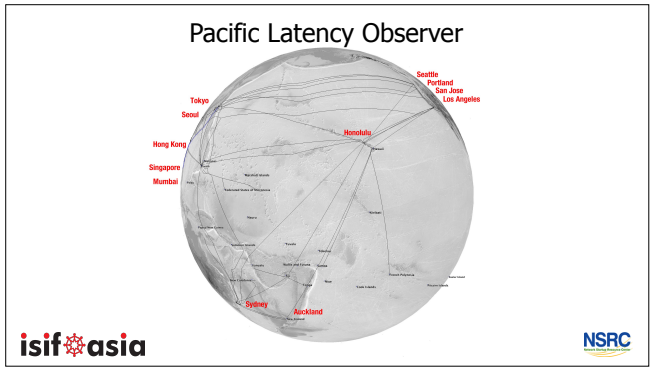

Theory is all well and good, but to test out these assumptions, I've set up servers around the world measuring latency into Pacific networks every five minutes. There they are, dotted around and on the globe.

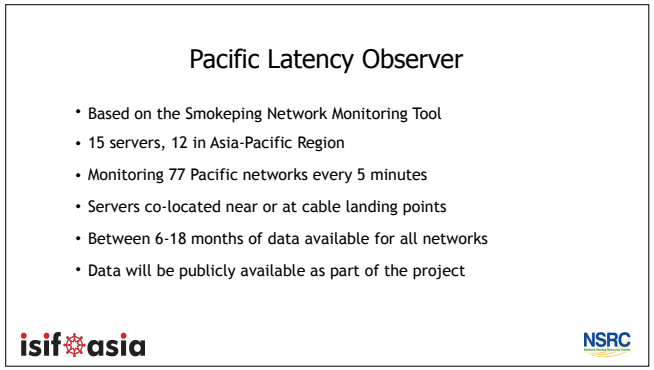

These servers are running a basic performance measurement tool called Smokeping, which tracks latency, loss, and jitter using ICMP ping messages.

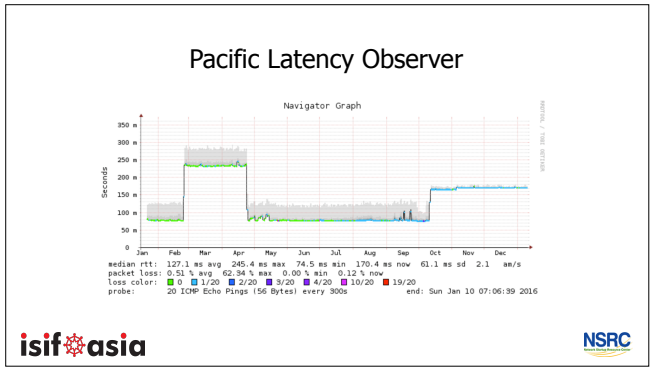

Results from the testing follow this form. In the chart above, we see that traffic between a server at DRFortress and Guam's incumbent telco followed a low latency Guam to Hawaii route for around seven months last year. At the moment, it appears traffic routes via California - we'll discuss how that can be confirmed in a few more slides.

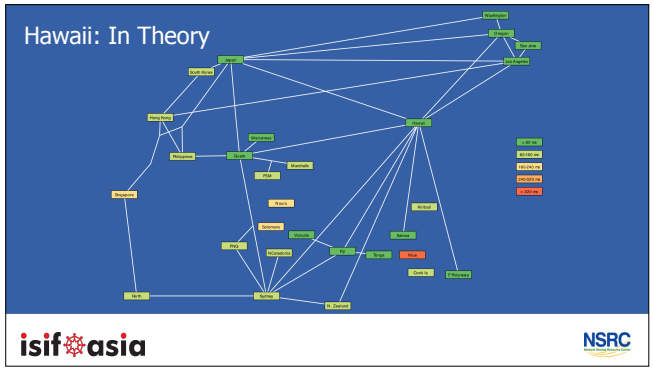

Flattening the world out, let's have a look at the Pacific from the perspective of traffic originating in Honolulu. With 03b connecting the Cooks, Kiribati, Nauru, & the Solomon Islands, we should be able to get to most of the Pacific pretty fast.

Why the colours? Remember back to our slide on latency. Most people will perceive delay when a communication takes 250 milliseconds. Above created 80ms colour bands. Less than 80ms, dark green. 80-160, light green, 160-240, yellow - caution!, Over 240ms, orange. More than 320ms, red!

A User in Hawaii should be able to Skype or Facetime with anyone in a major Pacific centre, with the exception of Niue, without perceiving any delay.

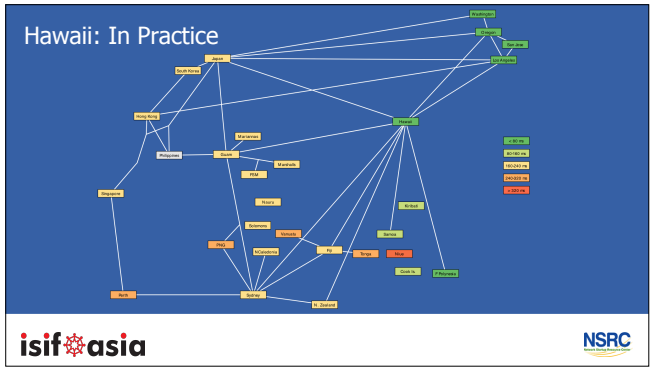

The reality is that there's a lot less green on this diagram than we'd expect. Why is that?

In practice, Hawaiian IP traffic reaches most Pacific destinations via the West Coast of the US, with the exception of French Polynesia, Samoa, and the Cook Islands. This includes Guam and Fiji, the Pacific's two other major submarine fibre hubs.

Physically Hawaii may be a hub for cables, but from the perspective of the Internet it's not the center of anything.

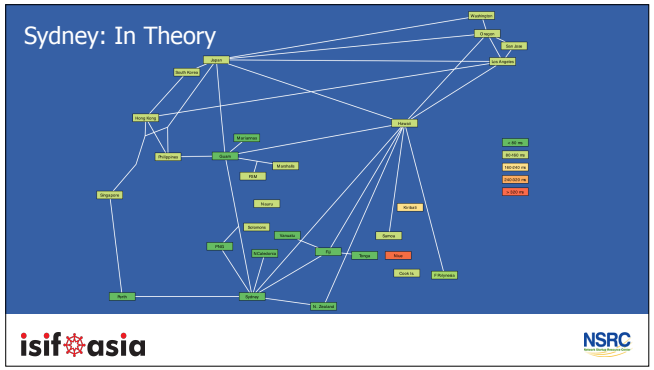

Let’s have a look at Sydney now. It’s well-connected to the Pacific via fibre and has 03b earth stations in Dubbo and Perth.

Again we expect a user in Sydney should be able to Skype or Facetime with anyone in a major Pacific centre, with the exception of Niue, without perceiving any delay.

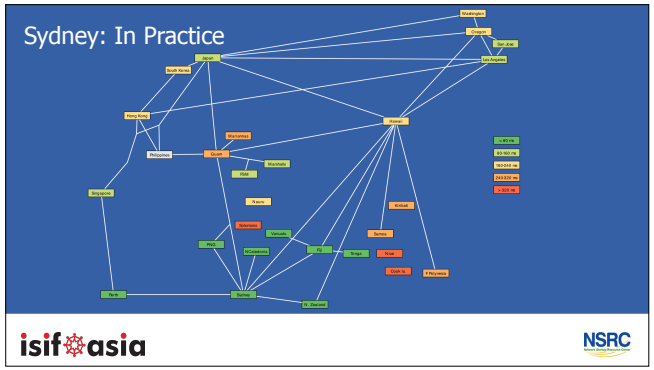

Practice is nothing like theory. We've already noted that Hawaii-Sydney traffic routes via the West Coast of the US. For Australian users, that means French Polynesia, Samoa, and all 03b subscribers connected to their Hawaii earth station are also reached via the West Coast of the US - including the Solomon Islands.

Reaching networks in Guam and the Marianas is also a trip through the West Coast, however astonishingly the Marshall Islands and Federated States of Micronesia are reachable via Tokyo - never mind the fact it takes two trips through Guam to get to them from Australia.

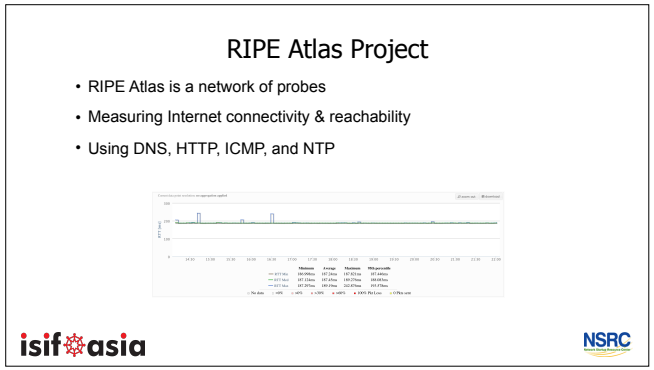

How could I possible know that traffic from Australian carriers to Micronesia and the Marshall Islands is routing through Tokyo? Another part of my project has been to work as an ambassador for the RIPE Atlas project. Atlas…

The devices look like this: They're tiny, they take very little electricity, they're secure, and they're useful. I've got one in my hand here.

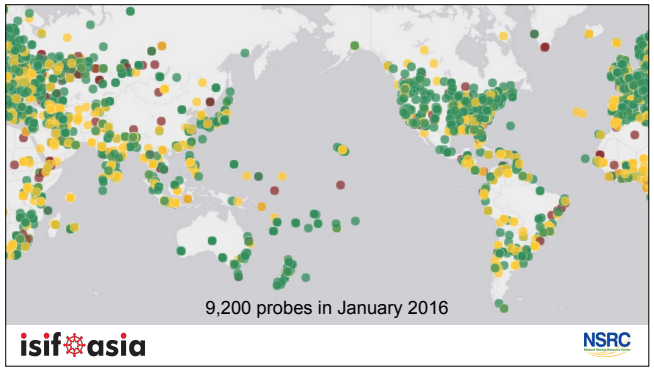

Here's a view of the RIPE Atlas network as of January 2016. 9,200 active probes.

Many of those dots in the Pacific are probes I've distributed as a part of my project this year. American Samoa, Australia, Cook Islands, Federated States of Micronesia Fiji, French Polynesia, Guam, New Caledonia, New Zealand, Niue, Samoa, and Tonga

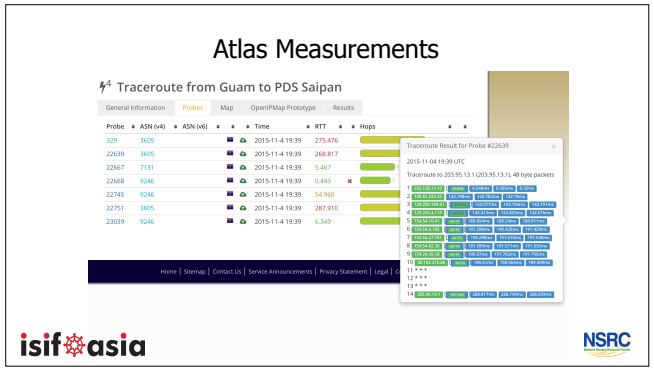

Some of the types of measurements I'm doing with Atlas are traceroutes. Here I've run a test from seven hosts in Guam on three different carriers to another carrier on Saipan, in the Northern Marianas.

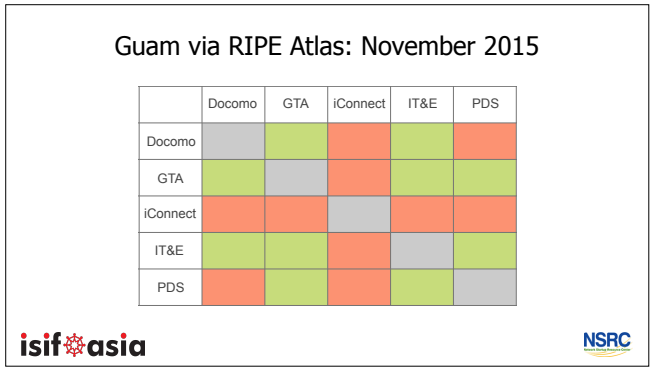

Here's the results of another set of Atlas traceroutes between the various carriers in Guam. We can see that not all of Guam's carriers interconnect locally - and some, like iConnect, send all of their traffic directly to the United States, not peering any of it on-island.

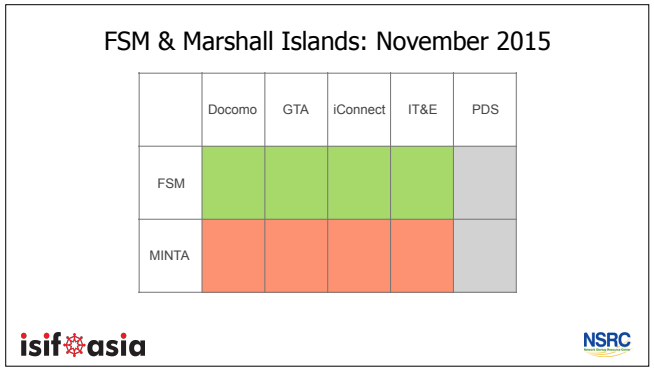

And here's the results of a set of traceroutes from Guam carriers to the Federated States of Micronesia and to the Marshall Islands. Now I showed this slide at PacNOG18 in Guam six weeks ago, and the CTO of Docomo Pacific took issue - he'd personally been involved in setting up the BGP sessions giving direct routing between the Marshall Islands and Guam carriers.

It turns out the peering session had fallen over and no one in any of the carriers had noticed. I'm sure a few users noticed - there's a fair amount of back and forth between Guam and the Marshalls, but the carriers, no. This has been fixed now, by the way.

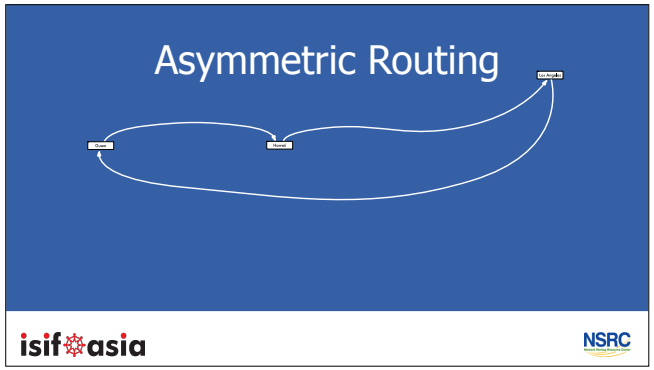

Another issue discovered via latency and traceroute observations is that of asymmetric routing. Also in Guam I made the assertion that none of the carriers had direct transit to Hawaii - and then I heard from an ISP called PDS that they did have a direct circuit, and in fact they were peering at the Honolulu Internet Exchange..

It appears from my measurements that while PDS have a direct link between Guam and Hawaii, they don't advertise all of their Guam routes at the exchange. From my server at the Honolulu Internet Exchange, the best route I can find for PDS is via Los Angeles. All of my traffic goes to California first, before going to PDS Guam. On the way back, PDS has learned a route for my server via the Honolulu Exchange and sends the traffic directly back to me.

This concept of asymmetric routing results in slower than necessary performance, and sometimes strange network problems. It's not confined to the Guam Hawaii path or PDS though. I see asymmetry with IT&E in Singapore, with GTA in Hong Kong, Seoul, and Tokyo.

Throughout my research I've been asked - especially by carriers - what are you going to do with all of this information? What's the purpose of your project? Is this all going to go in a report and never get seen again?

My intention with this project is to identify gaps in knowledge and produce a strategy for filling them. Roughly this help falls into categories described on the next few slides.

Help Stakeholders Understand: Network Visibility

- Where is your network traffic coming from and going to?
- Which networks are your “Top Talkers”?
- Are your customers being well served – and not just for Facebook?
  - Local peering is important for this
- Are you planning your capacity based on data?
  - Or just buying on salesperson recommendations?

Help Stakeholders Understand: Transit & Peering

- All transit is not equal
  - A link without committed latency can take a long, but inexpensive path
- Are you planning to satisfy today's requirements? Tomorrow's?
  - Long-term agreements must predict growth
- Do you understand peering strategies? The value of peer networks?
  - Free peering is great, paid arrangements are sometimes more appropriate

Help Stakeholders Understand: Streaming Media

- CDN content is available in the Pacific at Tokyo or Sydney
  - There's no reason to take CDN traffic from Los Angeles
  - Closer content is cheaper content
- Latency matters for CDN/Streaming Media Access
  - TCP rx windows restricted to improve CDN throughput
  - Distant users suffer to increase performance for all

Help Stakeholders Understand: RIPE Atlas Project

- Probes are free for networks - even multiple probes
- Assistance is available for many tasks beyond setup
  - Monitoring & Systems integration
  - Visibility from the world
  - Custom Measurements

Next Steps: A Peering Strategy for the Pacific Islands

- Integrated Pacific Performance Website Online
- Analyse Benefits of Regional Peering Points
  - Does every country need an exchange? Maybe not.
- Assess needs for training & assistance
  - Network Visibility, Transit & Peering, CDNs, Atlas

isif asia NSRC

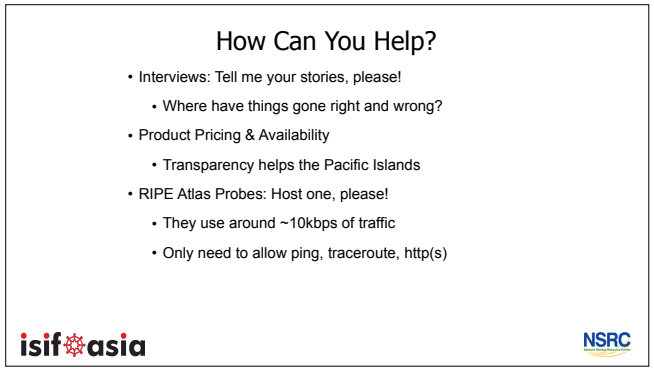

This is my wrap-up for now. In addition to handing out Atlas probes, I'll be interviewing users, carriers, and regulators in a number of countries to explore what's going on, and what we can do to help. If you'd like to get involved, please get in touch.


Thank You!
Email: jon@brewer.nz
Skype/Twitter: @kiwibrew
isif asia
NSRC