

IETF Technical Plenary Session
Monday, 21 July 2014
Toronto, Canada

Transcript provided by Brewer & Darrenougue
Corrections provided by Cindy Morgan and Andrew Sullivan

--BEGIN TRANSCRIPT--

>> RUSS HOUSLEY: We're going to start the meeting with three reports, a report from the IAB, a report from the IRTF chair, and a report from the RFC series editor. So after the reports, we'll have a technical topic, and then after that, we will have the open microphone time for the IAB.

So I want to go over a few highlights for things that have happened since we were last together in London.

The IAB sent some comments to the National Institute of Standards and Technology encouraging transparency and openness in all of their publications. This was in response to an announcement that they were going to be more open and transparent in developing the federal information processing standards but we thought it should apply to all of their security-related documents.

The IAB also commented to ICANN on the principles and mechanisms and processes that ought to be used in the development of a proposal to NTIA on the transition of the IANA function stewardship.

We appointed Sarah Banks and Robert Sparks to the RFC Series Oversight Committee.

We appointed Sean Turner to the Internet Society Board of Trustees.

And RFC 7288 on reflections on host firewalls was published. And then we appointed Matt Miller as the liaison manager for ECMA TC39.

We also got a workshop report on congestion control published, as RFC 7295.

We got a workshop report on the Internet Technology Adoption and Transition published as RFC 7305.

And we appointed myself and Lynn St. Amour to the coordination group on the transition of the IANA stewardship.

And lastly, we published a document describing the IEEE 802 and the IETF relationship as RFC 7241. That's actually an update to RFC 4441, or a replacement for it.

So this slide shows all of the documents that are currently being worked on by the IAB. I'm not going to go down each one, other than, say, two of them. The one at the top and the one at the bottom actually are intended to be BCPs, so the IESG will handle the last call and so on for those documents.

In terms of the standards process, we did not receive any appeals since London.

[Applause]

>> RUSS HOUSLEY: Ongoing appointments. We are trying to fill a seat on the 2015 ICANN nominating committee. We had asked for volunteers by last Friday, and we have extended that until this Friday, in the hopes that we can get a few more nominees or volunteers to become members of that nominating committee, so please contact us immediately if you're interested in that position. Right now, we have two people who have volunteered. We'd just like a bigger pool to choose from.

And the Independent Stream Editor, the IAB is seeking feedback on the job that Nevil [Brownlee] is doing. Basically, the feedback will determine what, if any, other actions are needed, because the contract with Nevil is about to expire and we have a choice of putting a contract with someone else or renewing that contract.

So a couple highlights from the IAB retreat: We met in conjunction with LACNIC in Cancun, Mexico. They had a meeting there, and so we joined them. And one of the big things we did is restructured the IAB programs to allow the IAB to focus on the topics that are most relevant to the Internet today.

There's a whole bunch of backup slides to this presentation if you want to download those from the proceedings page. There's a slide at the back about each of the programs and what we're going to focus on this year.

One highlight is that we previously had a Privacy program and a Security program. They have been merged to the Privacy and Security program, and it is going to be a focus of significant effort in 2014 for the IAB.

These are the current list of programs, and at the URL at the top, you

can learn about those. I don't intend to go through them. As I said, the backup slides about each one of those are there. And that's the end of my report.

Lars, you're next.

>> LARS EGGERT: Hello. I'm Lars Eggert. I chair the IRTF. This is the short version of the report. As usual, the long version of the report happens at the IRTF open meeting, which I think is Tuesday. It will be in the slides.

So we have four research groups that are meeting this week. That's a little bit fewer than normally. And we have two proposed research groups that are meeting.

The ones that are meeting are the Information-Centric Networking Group [ICNRG], Crypto Forum [CFRG], Software-Defined Networking [SDNRG], and Network Management [NMRG].

There's a proposed research group on Data Center Latency Control. They had a side meeting in London. They have an on-the-agenda meeting now on Friday morning.

And then there is a proposed new research group on Network Function Virtualization and they have an off-the-agenda side meeting, I think, during Wednesday lunch.

And there's a DCLC and an NFVRG mailing list at IRTF that also you can join if you're interested in those.

The IRTF open meeting, as I mentioned earlier, it is on Tuesday at 1420.

Unlike other IETF meetings, we are not reviewing the progress of one research group with the IAB this time. The reason being that fewer meetings so fewer chairs are here, and there also wasn't any pressing need because we reviewed most of them in the last year or two, so we decided to spend that time with the IAB on something else.

There's a third proposed research group that exists, called GAIA, which stands for Global Access to the Internet for All. They aren't meeting here, but they have a meeting in Cambridge in October and they have a meeting collocated with the ACM DEV conference in San Francisco in December, and the mailing list is gaia@irtf.org and if you're interested in that, join that list.

So my personal take of the research group's activity levels at the moment, you know, three of the four seem to be active at the moment. Three out of seven seem to be active. Four out of seven, maybe not so much. But that's not uncommon. Research groups go dormant for a number of reasons.

And we closed two since London. We closed the Network Complexity Research Group that had run a little bit out of steam, and we closed the Routing Research Group because that had run out of steam also.

That doesn't mean that we will never do any routing research again in the IRTF. It just means that now there's a clean playing field for somebody to come up with a group with a maybe new focus or new emphasis in routing and pitch it.

I believe it's better to close things and let a new group start rather than trying to keep something alive past its prime.

This is our stream. It's not very busy. We had three RFCs published since London, which is quite good for us. Two of them are from the DTNRG and one is from the Crypto Forum Research Group.

The DTNRG has published some bundle protocol-related RFCs and Crypto Forum accomplished the LTP encryption algorithm.

Finally, we are giving out the Applied Networking Research Prize. By now, everybody has probably seen this slide many times. This is together with ISOC, who is graciously funding this.

For 2014, we had 46 nominations of research papers that attempted to get selected for the prize. We have a selection committee that consists of a bunch of IETF people, some IRTF folks, some professors, some of the past prize winners we have invited. We picked six winners for 2014. Two of them have already presented in London. Two more were supposed to present here. One of them, Misbah Uddin, had trouble getting a visa to Canada, so he will present in Honolulu, together with the other two winners for 2014, and Robert Lychev will be presenting his research on the partial deployment of secure BGP during the IRTF open meeting on Tuesday.

We did actually add another talk to the IRTF open meeting because Misbah couldn't come and we had the slot, so we moved a talk on what you need to do in terms of having stable or reliable crypto in the age of quantum computing. We moved that because it's probably of broader interest to the IRTF open meeting, so if that is of interest to you, you have a second reason to come to the IRTF open meeting.

The nomination period for the 2015 award cycle is open now. I think it will close sometime in October, so if you come across a research paper that you see published or that maybe you published yourself recently that you think is eligible for the award because it's applied and it's relevant to us, please nominate yourself or please nominate, the author of that paper. We try to pick good papers, but we can only pick papers that have been actually nominated. So if you want to see good talks, help us find them.

Thank you.

>> HEATHER FLANAGAN: Hello, everybody. I'm Heather Flanagan, RFC Series Editor, and I want to take a moment to congratulate you all on an incredibly productive year. You've sent us a lot of documents.

This has been an unusually high surge of documents. At this point, we've already seen 70% of what we had last year in our system at this point in time. It's been pretty, pretty impressive. We were talking anecdotally and we can't remember seeing a bigger spike in the history of the series.

This means that the RPC, the RFC Production Center, has not met their SLA, which is getting 67% of documents that are in an RFC editor controlled state published in six weeks or less.

Probably the biggest factor to all of that, it wasn't just a question of a whole lot of documents; it's that several of those documents were sent in in clusters, and clusters are always a little more challenging to work with because there's a lot more cross-references and terminologies that has to be checked.

Some of these clusters might be familiar to you. I suspect some of you have worked on one or more of them.

STORM started our year with only five documents, but those five documents resulted in 574 pages. The average RFC is about 30 pages, give or take, so that was kind of the equivalent of 19 documents in a cluster all hitting us at the same time. Also this year, HTTPBIS, MANET, and TRILL, some of them had an average number of pages per document, but perhaps some more extended AUTH48 sections than we might usually see.

All of these clusters, by the way, are published at this point, but they were definitely a pretty big impact in the RPC's ability to actually get documents published in six weeks or less.

I have some graphs to help those who are more visually oriented see what kind of happened over the course of this year, and to be able to compare it to last year as well.

These are also to be found in the RFC Editor report -- links are included to that in the Ops and Admin review plenary on Wednesday.

One of the things I want you to take away from this is, you saw that really big column? Good job. But it's going down from there. It's trending in the right direction. So we're not concerned, but we wanted to make sure the community was aware that this is happening.

So summarizing the SLA, that's our chart that keeps rolling track of the last 12 months and where things stand.

Beyond that, we get into some of my hot and heavy projects, including RFC format. I think some of you may have heard that I'm working on RFC format and changing the canonical format from text to the XML will be the canonical format, and from there publication formats will be rendered. HTML, PDF, text, and ePub if we can get someone to help us get the requirements in. We'll do ePub eventually. I just don't know if it will roll out at the same time.

It's been very, very busy since our last meeting in London, and there's many drafts to look at. The tools development team is working on the statements of work that will help get these tools ready for when we actually roll the format out. And if you want more gory details and a chance to talk a little bit about the style sheet prototypes for what the HTML will look like, you'll want to be at the format session Tuesday after lunch.

One of the projects I've been working on fairly heavily since I started was the style guide and I'm happy to say that's in AUTH48 right now, RFC-to-be 7322.

Some of the significant changes -- I think probably the most popular change I heard from the community was the fact that the Internet-draft string will be included in the reference format. A lot of people really like that. Which is good, because it makes sense.

There's going to be a Web component of the style guide so that we have a section where we can develop guidance when new things come up or adapt things that change a little bit more quickly than we would want to publish a new RFC.

And you can see some of the other highlights there. That should be

published before too much longer. I'm waiting on my coauthor here somewhere.

Digital object identifiers. This is also something that's been very popular in the community. We are working on that. As folks know, the IAB approved starting the assignment of DOIs to RFCs. This will be to all published RFCs. Past, present, and future. It's not talking about Internet-drafts; just talking about RFCs.

And the tools development team and I are wrapping up the draft for the statement of work for the project, and you guys can expect to see that going out for community comment real soon now.

For the Internet-draft that's describing DOIs, we're leaving that in draft form right now because we'd like to get some real-world experience and then update the draft to match reality, and wouldn't that be a lovely thing.

So there's several ways to find out what I'm doing or to discuss some of the big pictures on format or suggestions on style or anything else you might want.

There's, of course, my wiki which I use as my own brain dump regarding what I'm working on, and I update it in surges as I can.

There's the RFC-interest mailing list. How many of you are subscribed to that? Cool. Gosh, you're like clumped together. How did you know?

And of course you're more than welcome to contact the RFC Series Oversight Committee [RSOC] to talk about any concerns you might have with the direction of the series or what I'm doing or anything along those lines.

And that's what I have for you. Thank you very much. Ah. David.

>> DAVID BLACK: David Black, STORM working group chair, the responsible perhaps guilty chair for that large cluster. That cluster, by the way, is the new version iSCSI specs which is a major component of the IETF and they're in widespread use and a credit to what we get done around here.

I wanted to thank you and commend the RFC Editor organization on what was done on those drafts.

You left out the rest of the story, which is, we had to ask you to put a bit of a rush on publication of those drafts for coordination of other

standards bodies. The RFC Editors set a schedule, kept to it. I was impressed with the professionalism, attention to detail, and patience of the folks doing it. This is one more example of why the RFC Editor contributes to the high quality of the IETF's standards. Thank you very much and thank you to all your organization.

[Applause]

>> HEATHER FLANAGAN: Thank you very much. Thank you on behalf of the editors who did so much of that work and continue to do so much of that work.

So thanks, all, very much.

>> RUSS HOUSLEY: Andrew, would you and your panel come on up?

>> ANDREW SULLIVAN: I'm Andrew Sullivan, and I'm your moderator today for a discussion about Internet topology and geography. Some while ago, Michael Richardson suggested to the IAB that we have a look at the IXmaps project Web site. He noticed a grant, I guess, that the Canadian Internet Registration Authority had given them, CIRA. And CIRA is the registry for .CA here in Canada. And, of course, we like to invite people who are part of the local networking community when we have these topics.

We were excited to begin with to talk to the IXmaps people, but also this led us naturally to think about some things that had been forced upon us anyway. So it was a happy coincidence.

There are, of course, people and governments who are keen to link intra-network traffic to geography, to geopolitical boundaries, and so on. But there are also real issues of geography that affect the way the network operates, and this is perhaps clearer than it would be in other places because, of course, Canada is huge, but has like practically nobody in it.

You end up with this funny geography and the network topology that goes with it. So today we have three speakers to talk about aspects of network topology and the way it interacts with geography.

First, we will have Antonio Gamba. He is a Ph.D. student at the Faculty of Information at the University of Toronto. He studies how the uses of do-it-yourself microelectronics constitute innovative forms for literacy as they entail new forms for collecting, interpreting, sharing, and negotiating information. But he's also a member of the aforementioned IXmaps project at U of T. He works with Professor Andrew

Clement who is there as well, and he is participating remotely because he is in British Columbia right now.

Jane Coffin is the director of development strategy at The Internet Society. She coordinates collaborative strategies for expanding Internet infrastructure, access and related capacities particular among emerging economies. And this includes ISOC's efforts to promote Internet exchanges.

Finally, Amogh Dhamdhere is a research scientist at the Center for Applied Internet Data Analysis. That's CAIDA. He does research in Internet measurement, Internet topology, and Internet economics.

Each of them will give a short little presentation about some aspects of their work and then we're going to have an open discussion. So I'm looking forward to the discussion. We already had a little mini discussion with them among the IAB. It was very productive, so I'm looking forward to this.

So without further ado, Antonio first.

>> ANTONIO GAMBA-BARI: Good afternoon, everyone. Thank you, Andrew, for that introduction. I'm Antonio Gamba-Bari. I'm part of the team of researchers at the University of Toronto.

As Andrew said, we are students at the University of Toronto. Our professor, Andrew Clement, is here and listening to us. So we are grateful at that IAB to have the opportunity to present our work to you in the company of ISOC and CAIDA on the topic of network topology and geography.

Since 2009, we have been developing the IXmaps Internet mapping tool, designed mainly to help individuals understand where their Internet traffic is routed and whether the National Security Agency may be intercepting their communication in their route.

We hope our work will raise personal privacy awareness as well as the informed public policy debates about Internet surveillance and routing.

We encourage people from disparate geographic locations and ISPs to install and run our traceroute program, to feed our database with traceroute through hostname parsing, latency comparison, and topological analysis. We geolocate the intermediate routers for mapping the routes their packets take. We highlight the exchange points where these routes pass through suspected sites of NSA interception.

As part of our public education mission, we have produced several videos to better visualize Internet structure and routing together with its personal and policy implication.

So I'm going to show one of these videos that we have created. So if you allow me...

---Begin Video Transcript---

>> By developing a traceroute generation program so that people can anonymously contribute their own routes to our database. We've also been reaching out to the general public by presenting our findings in lectures and informational events at the Ontario Science Center and by developing our explore page where users can check out geographical representations of the data we've collected.

>> Well, I hope we are dispelling some of the myths that get in the way of that effective understanding. In particular, the idea that the Internet is a cloud as an etherial space that doesn't have any borders. In fact, the Internet consists centrally of a few very important points, Internet exchange points and other switching centers. These are in big buildings in the center of cities. And who owns those and the deals that are made there are basically what drives the Internet routing.

>> Did you know that a lot of Internet traffic that begins in Canada and ends in Canada travels via the U.S.? We call this boomerang routing.

>> For instance, if I send an email to Nancy and she's at OCAD, which you can see right there, it basically goes down the street to the McClellan (phonetic) Physics Lab and the packets go from there to OCAD pretty directly.

If I send an email to Ontario Student Assistance Program, which is just over there, it goes from here to New York to Chicago and then back here. And I think that that sort of contrast is really a striking one and helps drive home the point about the physical aspects but also the peculiarities of routing.

We tend to assume that the Internet structure has a very efficient way of routing packets. But how the Internet is actually put together is through a series of private arrangements largely between carriers and for various reasons, they are very selective about who they exchange their traffic with.

And that's why there's such a difference between the two routes from

here to OCAD and to OSAP. It is because of the carriers that are involved. It is not about efficiency of routing.

>> One of the big problems with cyber surveillance, this tracking of our activity through the networks is that it's done invisibly. And it can be done on a mass basis all simultaneously.

>> Right now in the USA, the National Security Agency is conducting a comprehensive domestic surveillance operation. Since 9/11, in conjunction with the USA Patriot Act, the NSA has the power to surveil without a warrant the communications of U.S. citizens and anyone else using U.S. networks.

One such operation at the AT&T building at 18 Folsom Street in San Francisco was exposed by the whistleblower Mark Klein, a former AT&T technician.

He has provided documents showing that AT&T allowed the NSA to install splitters in its main communication links. These devices physically split the fiberoptic cable to produce two identical data streams. One stream perceived as usual, but its copy is fed into extremely powerful NARIS programs that are programmed to intercept particular types of traffic. Additionally, data intercepted by splitters is stored for later analysis.

Currently, the NSA is building a huge data storage facility in Bluffdale, Utah. With this facility, experts estimate that the NSA will have the ability to store a copy of the Internet and phone transactions of every American for decades to come.

Remember that almost all U.S. Internet traffic passes through these 18 cities. This means that the NSA can achieve extremely far-reaching and comprehensive surveillance of Internet activities by setting up their infrastructure in relatively few intercept points.

And because of the convoluted politics of routing, this is not only a concern for Americans. Canadians, too, are subject to this surveillance as much as Canadian Internet traffic passes through the USA.

---End Video Transcript---

>> ANTONIO GAMBA-BARI: Our work has shown that NSA interception in just 18 U.S. cities can capture nearly 100% of U.S. domestic traffic. Foreign traffic that transits the U.S. is also very likely to be intercepted.

From our data, we estimate that 25% of domestic Canadian traffic is routed via the U.S. and, hence, subject to NSA surveillance. We have focused on the comparative data, privacy transparency of carriers and ISPs.

We have shown that contrary to Canadian and European privacy law, they generally do a very poor job informing customers of how their Internet communication is handled and whether they are protected from state surveillance.

Encouraged by the response we have received so far, we are keen to transform IXmaps from a research prototype into a more widely usable Internet mapping and policy analysis tool. With funding from the Canadian Internet Registration Authority, we will be rebuilding our traceroute generation application to make it more reliable and flexible in the face of the challenges and changes in routing practices as well as make it easier for individuals to initiate their own traceroute generation.

We also need to automate and improve the accuracy of our geolocation in part through crowdsourcing and improve the accuracy of technical backgrounds in these areas which we are currently doing.

We don't have a strong technological background in these areas and so we are currently looking for experienced developers to work in our traceroute generation and geolocation practices -- geolocation modules.

So far we have focused on North American routing. While our techniques are applicable elsewhere, we don't have the resources to respond adequately to the interest we have received in other countries. We welcome offers of help internationalizing IXmaps and making it more sustainable. In particular, we will put it under a free Open Source software license to make it easier for others to take it in their own directions.

And, quickly, I wanted to point to our Web site, if there are any details that you want, we are welcome to answer them. By the end of the month, we will be releasing these calls for proposal for recruiting developers and other network analysts that could help us to push this project further. Thank you so much for your attention.

>> JANE COFFIN: Good evening. And thank you very much. My name is Jane Coffin. I work for The Internet Society. Although I'm standing up here by myself, think of about 15 to 50 other people who do work with us. We do work in partnership inside The Internet Society with colleagues around the world in different regions but also with

organizations like the regional Internet registries, EuroIX, other IX associations like AXIS, AfPIF and LACIX and others. This is to develop Internet infrastructure. And, again, I would agree, the cloud is not something we're working on. We're working on hard-core infrastructure, building local infrastructure.

So I have colleagues here in the room with us. If you have any questions, if you speak Spanish, Christian O'Flaherty is here and we have other ways that we can talk about what we are doing with other partners.

What is an Internet exchange point? Many of you do know. It is a physical location where different IP networks meet to exchange traffic, and keep local traffic local. This is important to the Snowden debate later, and I will come back to that at the end.

Often people will say, "Oh, Jane, they are just boxes and wires." Well, we know from past experience and experience with our partners that 80% of the engineering in an IX is human. 20 is generally the technical. Philip Smith and I were in Papua New Guinea and he said, nope, it is 95/5.

95% of this is the human engineering, how we bring the different character sets together of ISPs, network operators, research and education networks. They are a vital part of the Internet ecosystem. It is not a magic bullet. It is not the one and only thing, but it is critical to develop the human, the technical, and the governance infrastructures in countries. Again, keeping local traffic local is one of the aspects. You build local Internet communities of interest that may not have existed or strengthen those that were latent or just in development. You improve the quality of Internet services. You drive up demand. Latency comes down; quality of service usually goes up.

It is a convenient hub also for attracting and hosting infrastructure. Content is key, as many of you know. Content is generated by businesses that have confidence in those infrastructures. And we know this is a catalyst for overall Internet development from our experience and what we've seen.

We've put out reports to help document the impact of IXPs. My colleague Karen Rose was instrumental to this study as was Michuki Mwangi in Kenya and some of the IX types in Nigeria.

The important thing here that we were able to prove particularly in Kenya from 200 to 600 milliseconds, those of you know that know how important this is, reduced to 2 to 10 milliseconds. In Nigeria, from

200 to 400 milliseconds to 2 to 10 milliseconds.

There were direct savings for some of the mobile operators on international transit. This was important particularly in developing economies.

These IXs also help facilitate eGovernment services. In Kenya, the tax authorities peer at the IX. It was critical for the local infrastructure for them to peer at the local infrastructure. They also had to catalyze local hosting and content, as I said, increased the mobile data market and attracted regional traffic particularly to KIXP. So their international networks are coming in to peer. So it comes from the national to the regional and now they are attracting international peers. We have also put together a study, again some data for you, to prove that in Latin America, some of the findings we've seen working with the professor at the University of San Andreas in Argentina in one city alone after the IX was installed, they went from \$100 per megabit per second down to 40. This is huge.

Recently when I met one of the key people, Antonio Harris, at CABASE, which is the network in Argentina, allowed IXPs. He told me just the threat of CABASE coming into a city recently brought down the price per megabit per second from about 250 to about 10 or 20. It is really interesting to see what happens in these environments where there was some competition and it is growing, particularly when the IXs come in.

In Brazil, at NIC.BR and the PTT Metro System, there were 26 IXPs, and we recently chatted with them today.

They are up to 600 gigabits per second at peak. That's huge. They've created an amazing system working with the government and companies across the country.

In Ecuador, before the IXP went in, international transit was \$100 per megabit per second. It is now local traffic. Local traffic, local is \$1 per megabit per second.

They are also running RPKI. You will hear more about tomorrow, I think at the OPSEC meeting. It's at 1300. Give it a go to have a look because it is interesting to talk to the people that are deploying new technologies through the IXPs.

IXs also can be a place to host your ccTLD, put in a time server, and deploy IPv6 as well.

After the CDN cache was installed in Quito in 2009, traffic went up by

700%. This is local traffic.

Additional work that we're doing in the region -- and, again, this is inside the region with the people that speak the language and know the regional actors -- putting in 12 Raspberry Pi's in Bolivia to take a look pre IXP and then to take a look during the IXP traffic deployment and whether the service and the quality of service goes up or down. We'll take a look there.

Network efficiency, studied in Argentina with a professor at the University of Buenos Aires again to take a look at CABASE's network to see how network efficiencies can be improved and how we can help them there.

The study can be found at the link in the text here.

They are over 350 IXPs around the world. There are different mapping technologies being used to take a look at that. This is a telegeography map. There are also maps on our Web site, at the ixptoolkit.org. There is one at euroix.net, and Packet Clearing House also has mapping.

When we take a look at building capacity in differing countries, we often will, of course, consult with the local community, work with other local actors, try and take a look at how we can make the IXP grow if it exists, how it can be a part of that ecosystem, and what the right business models are. Usually it is leveling up over time from a startup to an entity that might be charging for some of its services. We often help with equipment as well because the startups are difficult.

With technical skills, we're taking a look at helping them with routing, network management, network efficiencies, running the IX and working with the local Internet community and different authorities.

In a lot of countries, you do need to work with the government. That may be antithetical for some, but it's very important to figure out what the enabling environment can be for the government, working with the technical community that you're building.

In Africa, we have two projects called the Axis I and II. This is with the African Union.

The Axis I project was something we bid on and won due to the great guys that we work with in the region. There were 30 best practices workshops, bringing the community of interest together to kind of coalesce. It's not easy.

And then 30 technical aspects workshops. These are a bit more hard-core. This is BGP, how to run it, engineer it, how to interface with the IX and also whether or not you're running v6 and other relevant services.

There have been four IXPs launched this year, and this is quite something, because it started in 2012. In 2014, we've been working with AfriNIC, Jaguar Networks out of France, Lyon-IX, and INEX, which is the Internet neutral exchange in Dublin.

Axis II, five regional meetings across Africa to bring regulators, policymakers, network types together to explain the importance of growing that regional infrastructure. There are so many landlocked countries in Africa that it's important for some of those government entities to try and work together.

There was one instance in Zimbabwe where it took almost two months to string some fiber about a hundred meters, due to the fact that it was over a bridge that was historic.

There were different authorities from the customs officials to the border officials to the communications ministry, and they've often come to us and others and said, "Can you do something? Talk to some of these officials. See if you can help us facilitate a discussion. Because how is it that it takes, from a business perspective, over a month to deploy that fiber? It's not good for business. It's not good for the community. And it's not going for deploying that infrastructure."

We also host something called the AfPIF in Africa, the Africa Peering and Interconnection Forum. The next one, the fifth, which is quite something, it's kind of a benchmark for us, will be at the end of August.

This brings peering coordinators together, network managers together. It's a really interesting event where IXPs that are starting out can talk to more robust IXPs. They also have a chance to interface with each other, have meet-ups with some of the network operators in the region as well.

Latin America. We knew that, looking at certain countries that had deployed IXPs, certain factors existed from the market conditions and the regulatory and policy environments. Those are Argentina and Brazil. That's over 16 years of IX deployment. Countries that we're working with now who have recently deployed IXPs had some of those conditions that weren't as strong. Not a very strong Internet community. Some of the market conditions, again, from the policy and regulatory side were

down. It's why we work so closely with governments. We provide some pre-best practices training where we invite companies in through the government. Sometimes it's the only way to get them to the table.

We have a joint training objective, which is also training the government to appreciate that the technical community should run the infrastructure.

The government can help enable.

We've seen some interesting cases in Costa Rica and Bolivia, where in Costa Rica the government was less restrictive and the environment has progressed faster for developing the IXP and I believe they just launched it.

In Bolivia, it's taken us a bit longer. The IX is now embedded in the law. It must be licensed. Anything related to interconnection is also being heavily regulated.

So we work as much as we can, slowly, slowly, it's step by step, with the government, with the community, and we try and build up that Internet community as we're going so it can do what it can with the government. It's not an easy process.

Again, some of the training we do do in the region, introduction to BGP and traffic engineering. We do a lot of joint training -- and I want to stress the partnerships here: LACNIC, LACNOG, PCH, governments, company experts -- in the basics of architecture and how to obtain resources. You'd be surprised at how many operators don't have ASNs and may not know where to get some of those addresses. They're working through their up-streams. So they will go into Africa and other Latin American countries. They'll talk to LACNIC or AfriNIC about obtaining those resources. But it still happens.

Some people don't have ASNs in their infrastructure or they're not -- don't have the right addressing plans as well.

With respect to equipment, I mentioned earlier the start-up fees is often one of the most difficult. We help donate new equipment to IXPs to help them get started and help install that equipment and explain to them how to keep it going.

We do help level up existing IXPs, which means that if they need a new switch, router server, they can come to us and we often can help them with that donation.

Partnerships that have been developed in that region, development of

LAC-IX was through the Internet Society and LACNIC. ISOC has been working with LACNIC to train partners in the region. The RPKI deployment was one of the key things in Ecuador and that was thanks to LACNIC.

Community building. Regional Interconnection Forum. This is within the LACNIC meeting. There is a LAC Peering Forum. This is a working group within LACNOG.

I didn't stress earlier, but the network operator groups are critical to that infrastructure development and Internet technical capacity training for people in the region.

Some of our other partners there again are LACNIC, LAC-IX, NIC.BR, Packet Clearinghouse, LACTLD.

Working with governments, critical. Companies and organizations like Cisco and the Google Foundation. I'd be remiss if I didn't mention them as they provide grants and equipment to us.

We also have an IXP toolkit and best practices project grant. This is more of a global project to look at places where there isn't as much of a robust infrastructure or we can help facilitate other partnerships.

The grant builds on our previous efforts, and it's creating an IXP toolkit report. It's a study and a methodology to identify best practices. Creating and improving an IXP portal. Fancy name for a Web site.

Partnering to conduct training and hold workshops, building capacity around the world. And that's working with academics, the IX associations, other IXPs who have great experience around the world, and the RIRs. And NSRC in the works in Asia-Pacific.

We've worked across eastern Europe. Recently in Papua New Guinea someone was interested.

We don't often have to be the end entity that's in place in the region to help out. We can often hand off. So we're not ever trying to be the one and only, but bring partners together.

As I said before, Philip Smith, who is with us, many of you know him, he's got amazing experience, and NSRC is very interested in making sure they can maximize some of the work in the region. The last thing I'll note, in the toolkit we're using open source software for our mapping technology. It's a little wonky at times, but it's kind of cool when

you look at this because there are a lot of different IXs that want to see themselves on the maps, so it's something of interest to them. We got about 25 hits at the APRICOT meeting in Malaysia in February where we rolled out the toolkit.

Now, Snowden-related aspects of this. Keeping local traffic local is a critical part of that IX infrastructure.

If you are tromboning your traffic or boomeranging it, if you're sending your traffic from Nairobi to London to exchange between two ISPs in Nairobi, that's not efficient, right? Time and distance equals money. A lot of other aspects there. But if you're exchanging traffic locally with each other, you're driving down costs, you're increasing your quality of service, and you're driving down the latency. This is a critical thing.

We do get questions from local governments now about the Snowden impact on IXPs.

What is the IX set up to do? Well, obviously it's to peer and switch exchange traffic. It is not set up to be a monitoring facility for deep packet inspection. Or at least that's our philosophy.

Could someone do that? Yes. But do we come in with that? No. Just wanted to be clear. This is a question we get very frequently now more often.

Again, we're trying to keep local traffic local for network resiliency reasons, for development of that Internet infrastructure, for local human capacity development as well, and for content creation and design because that's going to develop that infrastructure more as you go forward.

In any event, I'll leave you with all that, but this is part of what we're doing around the world with partners, and again, thank you very much for your time.

>> AMOGH DHAMDHERE: Okay. Good evening.

My name is Amogh Dhamdhere. I'm a researcher at CAIDA, and today I'm going to be talking about some of the tools and the data sets and research that we're making available to the community to support sort of the general theme of topology and geography research and analysis.

So CAIDA is Center for Applied Internet Data Analysis, a recent name change, actually, but we kept the same letters.

We're a research group at the University of California San Diego. We are at the Supercomputer Center. We do look at both practical and theoretical aspects of the Internet. We do a lot of network measurements and data collection and data sharing.

So jumping quickly right into the theme of today, which is topology and geography, we have built and operate an active measurement infrastructure which is called Archipelago, or Ark, and as of today, it consists of 102 monitors and it grows pretty consistently at about one or two a month.

37 of these monitors are IPv6 capable. 39 countries and 88 cities are represented. And 54 of these monitors are actually little raspberry pi computers, and those are pretty much what we hand out these days as our monitors.

Lots of different projects run on this infrastructure. We have ongoing team-probing experiments to collect IPv4 and IPv6 topology by doing large numbers of traceroutes to pretty much the entire routed v4 and v6 space. We do alias resolution measurements to collect router-level topologies. Recently, we've started using these monitors to do some measurements of inter-domain congestion to figure out where congestion is on peering points between networks. And various other research measurements such as spoofing measurements are also run on these monitors.

So as I said, just on the last slide, the monitors that we hand out these days are pretty much these Raspberry Pi computers. We all know what these are. They're not super-high-power devices but they're good enough to do the job that we -- that we want from them, which is basic network measurements.

The good thing is they only cost \$35 apiece and it's pretty easy to hand them out to people and get them to host them.

So we're always looking for new vantage points, and if you're willing to host one of these either at your organization or at your home, we would be most happy to have another vantage point.

So the Archipelago infrastructure has actually been operational since 2007 when the IPv4 topology probing started, and IPv6 started a year later in 2008.

It's generated a huge amount of data coming out of 6 terabytes of compressed data at this point. This consists of both IPv4 and IPv6.

You can see here on this picture the number of monitors that we've had over the years, and it's steadily growing. This is actually a picture from last year, so there are many more at this point.

In terms of the data that we make available from the Archipelago infrastructure, we have the raw traceroute data. This consists of IPv4 and IPv6 traceroutes. This is approaching 6 terabytes of data at this point and this is available in its original raw traceroute form, if you want it.

In addition, we make available curated topology data sets, which we call Internet Topology Data Kits or ITDKs. We do about two of these a year.

And this consists of a router-level topology, a router-to-AS assignment data set, DNS names, geolocation, and other fun stuff.

We also provide traceroute-derived IPv4 and IPv6 AS topologies.

So there is a lot of this kind of data available and it's historical, but one of the goals of a currently funded project that we're working on is to make it easier for researchers and people interested in this kind of analysis to actually access this data and do interesting things with it.

So we're building support for rich queries on this traceroute data, and the idea is to put them together with other kinds of data such as geolocation, annotated AS-level topologies and router level topologies, and router-level topologies.

So for instance, just to throw some examples out there, if we wanted to look at all the traces from a monitor in Canada -- and we have two of them at this point -- to destinations in Canada and traces that actually exited Canada and entered the United States through a few hubs and came back in, then we would like to be able to support this kind of analysis.

Another example, suppose we predicted that a certain region was going to be affected in the sense of a natural disaster like a hurricane or a storm coming up or political instability. We'd like to know all the paths from our current monitors that actually traverse this region. These paths might be rerouted or might even go down when something actually happens.

So I could go on with these kind of examples, but this is very much a work in progress, so one thing that will be very nice is to hear from

the community about what kinds of analysis you would like to do with this kind of data, and if it can be done with the data, we'll be very happy to support it.

Maybe by the next IETF meeting we have something to show you.

Ark provides an interactive topology-on-demand service, so this is called Vela, and it's essentially an interactive interface to all the Ark monitors that we have, so once you get an account on the system, you can log in, configure some basic measurements from the system, which are currently pings and traceroutes. You can choose to do these from a single Ark monitor or you can choose to do them from multiple Ark monitors, and there are various forms of the visualization output. The one I've shown here on the slide is a geomap which shows the traceroute, basically geolocates each hub of the traceroute and shows it to you on the map.

In addition, we've recently built a system for DNS-based geolocation and this consists of two pieces called DRoP and DDec.

So DRoP is actually an algorithm that automatically picks up DNS -- geographical hints from DNS names. It uses data that we collect as part of the ITDK, which is RTTs and TTLs towards various interfaces that we've seen. It then uses all of this data to build geographic hints, validate them, and then actually it also tries to extract general rules for a certain domain in the sense of how interfaces of that domain seem to be named.

We have a public-facing component of this system, which is DDec, and you can access it actually at ddec.caida.org. It allows you to enter a list of DNS names and tries to infer that you -- if any geographic hints are present in the DNS names, it tries to return to you the geolocation, basically.

The public interface also supports correction, so if you're -- if you're operating a network and you have a certain generalizable rule that you can feed into our system, that helps the whole process.

As-rank.caida.org is a repository of tools and data that have to do with autonomous systems, so we have business relationships between autonomous systems, customer cones, and ranking of ASs.

In addition, we provide for each AS a PoP-level map which uses geolocation data to infer where this AS seems to have deployed some infrastructure.

We also provide a router-level map basically using our Internet topology data kit and putting the routers that we infer for the AS on a geomap.

Again, as with the previous case, we support feedback from operators so we have an interface where operators can go in and enter corrections to the inferences that we've made.

Finally, very quickly, pointing to some pieces of recent research coming out of CAIDA, so last year we had some work trying to infer which networks peer at which IXPs using the route servers, and we've developed a method to infer multilateral peering at IXPs, and this is still a work in progress. We're trying to expand the site of Internet exchange points from which we can actually infer reliably the set of connected networks.

We've recently done some work on mining historical peeringDB data so this is data from the peeringDB that we've been collecting historically, and we've mined it to figure out collocation by different networks at IXPs, what kind of peering policies they advertise, how all of this evolves over time, and we can actually find interesting things like geographical expansion of network just by looking at historical peeringDB data.

We're currently in the process of investigating connectivity in the LACNIC region. This is very much a work in progress, but the idea is to look at sort of connectivity in that region at IXPs and how paths are coming in and out of that region.

A couple of years ago, a team at CAIDA was looking at the analysis of country-level Internet blackouts and outages such as those that happened in the Arab Spring when some countries just basically switched off connectivity to the rest of the Internet for their users. And they did this using a variety of data sources like BGP, Ark traceroutes, and also using data from the UCSD network telescope.

The same team also did a follow-up study looking at the impact of natural disasters, such as earthquakes and hurricanes, on Internet connectivity in different regions and also trying to develop metrics and tools to automatically detect outages of this type.

So all of these publications are on line. The data that went into this research is mostly available to the community, and so are most of the tools and techniques that we've developed, so a lot of this can be found on our Web site and if you'd like to collaborate on anything or just get access to the data, then we'd love to hear from you.

Thanks. That's about it. And I should mention at this point that a lot of the work that I presented was by many different people at CAIDA and it was spanning several years of work, actually. Thank you.

[Applause]

>> ANDREW SULLIVAN: So thank you, all. I appreciate the time that you've taken to prepare this and to tell us about some of the work that you've been doing.

I just wanted to start the conversation, and we're going to open the mics in a moment, so if you have pressing things you want to do, you know how to get to the mic.

But I wanted to start by asking sort of a "so what" question.

So we started with some evidence about boomerang routing and there's the suggestion that we want to keep traffic local and so on, but on the other hand -- right? -- people are paying money for this, and if they really wanted it to be better, they really wanted the traffic to stay locally and so on, wouldn't they just spend the money on it?

So I don't know who wants to hit back, but... I don't know. Is it working? Give it a try. Just talk.

>> JANE COFFIN: In certain countries, they can't afford the infrastructure. Chad, for example. Or Cote d'Ivoire, where there is a very strong incumbent and you often have to go in and try and build that community of interest around a very strongly financed entity.

So I'm not going to say it's sort of the undercover technical community we have going in, but sometimes you go in and build up that community. Then you work with the government to try and figure out how to better interface with them to encourage them to understand the importance of that IX, to build out that market, and it may not actually seem too obvious at first. But when we help with some of that start-up, if you will, it's important to bring in that equipment.

We're not talking about used equipment, because some countries are a little twitchy about that (as are the export control authorities) but you do bring in that new equipment, you help them install it, but you're not leaving people by themselves to not understand the technology building up that community and building up their own technical capacity.

So that will build up a stronger infrastructure, Internet

infrastructure, in those regions.

I would just say that aren't a lot of companies who are going to go in there on their own. That's why we work with the Internet technical community, the RIRs, NSRC, PCH, and others who are good at doing this. Many of you in the audience, maybe. We often look for volunteers, if anyone is interested in training and helping as well.

>> ANDREW SULLIVAN: Maybe the IX project, though, has some insight on this. Why does the Canadian traffic all go that way? I think it is pretty obvious if you're going to go from Toronto to Vancouver, the shortest path is going to be through Chicago. But pretty clearly the shortest path between two buildings in Toronto is not via New York and Chicago. So what is the financial interest there? I mean, I think I know, but I guess I'm going to ask you anyway.

I guess one of the questions that I really have about this is given that all of this data that you have collected, is it possible that we can tease out of that those kinds of relationships so that we can see, wait a minute, these things are happening this way and there's some sort of negative network effect? It is essentially a network economic effect that's at work.

So I don't know. Antonio, maybe you have some idea about why this is happening.

>> ANTONIO GAMBA-BARI: Okay. I think this is one very interesting question that you are asking. From our research, we have been sort of finding these patterns. What we call the boomerang routing has become more and more interesting for us in terms of questions about privacy, in terms of questions about Internet sovereignty. And later on, the more we get into these type of questions, we understand that on the one hand, there are reasons why the boomerang routing or the cross-border routing has very specific reasons why it happens. As Andrew just mentioned, there is no capacity to go from Halifax to Vancouver, or there are agreements, or actually there are ready networks that are in place that it would be more efficient to go in that direction.

However, some of our findings -- and that's one of the videos -- one of the traceroutings that you saw in the video, that we run a traceroute test consistently from the University of Toronto to the OSAP, the Ontario Student Assistant Program. And what we discovered is some networks decide to peer with some and not with others.

I know Andrew Clement, my professor, is listening to me right now. So I wonder if he wants to ask something here because this is a very

sensitive subject because it raises questions about not just efficiency but questions about whether our data crosses some borders and because of that reason is subjected to other inspections or potentially is subjected to data inspection.

So I guess that's what I will say.

>> ANDREW SULLIVAN: But I think -- and I'm sure there are people in the room who will observe this -- that, right, the routing isn't going to protect you because we know that some of that surveillance goes on inside the Canadian border as well. So that isn't really protection.

So then fundamentally, I guess, the question is: Does it make the network worse if you have that? And, I mean, it seems to me it does. But I guess what I'm wondering is: How much worse is it? What do the numbers show us? If we look at the data you've got, is this really a big tax? Probably in Chad it is. But how bad is it if some packets have to go to New York first in order to get across town? Do we know what that costs? Presumably it is real money.

>> ANTONIO GAMBA-BARI: I'll jump in -- we've had this question before. Very interestingly, some representative from TORIX, which is one of the main Internet exchange points in Canada -- as far as I understand and as far as we understand, there are deals between big Internet -- big ISPs and small ISPs. So the question about whether it's more expensive, I presume it has to do with who owns the network and what kind of deals are taking place in those exchange points and to what extent the people or the partners that you decide to peer with are mediated by the economic or financial reasons as well.

So the idea of the cost I think is related with the peering, the type of agreements that are taking place there, and, of course, in terms of relaying how long it has to go. You were mentioning before how increasing or decreasing the latency is actually some things that we have to do and we are aiming for. But we know less of that, and we are sort of learning that this is something that has to be a bit careful.

>> AMOGH DHAMDHARE: One thing I would just add to that is perhaps the role of longitudinal measurements in trying to tease apart some of these issues.

I mean, there is data that exists about business relationships and characterizing whether two networks a customer provides appears, and these things change over time.

And we can track these things historically and perhaps correlate them

with performance data in the form of the sort longitudinal traceroute data that we have.

About the issues, okay, this peer changes at this point and how did it affect a latency going from monitor to certain destinations? And this might shed some light on the question of what you were getting at with respect of how much these things actually cost in terms of how much better could it be.

>> ANDREW SULLIVAN: Do you think that the dataset that CAIDA has is at least a piece of that, or do you think there is data sitting there that people could tease that out?

>> AMOGH DHAMDHERE: I do think so, yes. I mean, there is a lot of traceroute data, like I said. And we also have the AS relationships and AS level topology data historically. So perhaps, you know, putting these two together can give some answers.

>>ANDREW SULLIVAN: I see someone at the mic.

>> JACQUES LATOUR: Hi. Jacques Latour. I'm with CIRA. We are one of the sponsors for IXmaps.

At CIRA, we are .CA so we have been doing research in Canada for IXP in the last, I guess, three years. We had two IXPs in Canada, and now we have close to seven.

And so one of the big changes that we've seen with putting in a new IXP like in Winnipeg is, first of all, you have to build a community to get the IXP together. And what changed there is the cost for bandwidth went way down. So the fundamental problem is the incumbent in Canada. They don't want to peer with local ISPs because they want to bring the traffic in the state and make money there.

So right now we're building what we call a donut architecture in Canada around the incumbent. And the challenge is getting them to peer in Canada with IXP. I think it is legacy. They think that if they are going to peer, they're going to lose revenue. And that's pretty much where we're at.

So the big thing that's happening now because of all the new IXPs is tier one carriers are moving in Canada and now competing with them directly on their soil.

So now that we've got new IXPs, the landscape is changing; and I'm not sure where all of that is going to go. But price is going down right

now for bandwidth. So it is all good.

>> ANDREW SULLIVAN: So, Jacques, before you go, I alluded at the beginning of my opening remarks that, you know, Canada is a big place. There is not very many people in it. And there's a lot of sparsely populated areas, of course, that presumably are not getting these IXs. Or is that one of the targets of your program?

Are you explicitly going after places that are perhaps underserved because there's nobody there?

>> JACQUES LATOUR: Well, the idea is to go coast to coast. The core of the Internet, the top of the Internet is the IXP. That's where you generate bandwidth. That's where content providers go. This is where people can get high volume of data for low cost.

So Halifax is now up and running, which affects IX. They have got ten peers. New Brunswick, we are working. So pretty much every province all the way to Vancouver we have new IXPs running.

The key thing is that we need to work with them to build a community so they can be sustained and operational on their own. We can't run it for them. We can't help them. It is growing. Even Kingston is looking at building an IXP.

>> ANDREW SULLIVAN: Thanks.

So, Jane, that sounds like something that's fairly similar to your experience in some ways. But, of course, you're dealing in areas where you've got lots of population but you've got maybe other incentives that aren't working.

So you were mentioning that there is some concern about people moving in because they come in with the IXP and maybe it becomes a point of surveillance or whatever. But this sounds to me like a sort of false problem, right? Because the monitoring is already happening.

So then question is, can we measure out of these things the positive benefits in those localities, some of the numbers you were showing were pretty great.

And that's interesting, of course, is what Jacques just said. It is going down here, too. Is this a widely spread dataset? Is it shareable? Can people find that stuff? Because maybe that will encourage this development.

>> JANE COFFIN: The datasets that we do have for Nigeria and Kenya are in the report. The link is in the presentation.

We also have data that was newly collected through the researcher in Argentina and the Latin America-Caribbean report. We are working very hard to try and collect more data. It is not easy if the IXs don't have the software installed already. And that's something we are trying to do in different parts of Africa and with existing IXPs that have been developed.

Put the servers in, work with them to understand that it's not surveillance software because that's a question that is often asked, and encourage them to also put up the statistics for their daily, weekly, monthly traffic and annual traffic. If companies don't see that their eyeball is on the network, it is difficult to think about why you would go in and invest for some companies. Others get it. They're already doing work in those regions. Like, Lyon-IX from France is actually working very closely with French-speaking West Africa to help us to do some of the training; Jaguar Networks, because they see this is the next community of interest that is being developed and those infrastructures will grow eventually.

But as Jacques said, it is very difficult to grow those communities of interest just off the bat. They are not mushrooms. You don't just sprinkle some water and they come up. It takes a lot of time and energy, two to three years. In Trinidad, it was seven. They just launched, I think, last month.

I know from my colleague Christian, it is two to three years right now in Bolivia. And it will take some more time because you are growing that community of interest.

And you want it to be sustainable so it is not something that you just impose. It is like what many of you did so many years ago. Just think about where you started. People are starting over in many places.

So it is important to bring that along slowly and not to try and force anything so it doesn't look like we're trying to do something that would have surveillance implications.

There are governments who are worried about this, and that's where we do get those questions, a lot of training workshops.

Again, a lot of those governments aren't monitoring the networks per se because there is so little they monitor. But they are probably thinking of doing that. It is just a fact. So the question is: What do we do

to encourage a proper technical infrastructure that's devoid of that issue?

>> TIM SHEPHERD: Hi. My name is Tim Shepherd. [For] Antonio, just a fairly simple question. You were showing the path between the two routers that showed up in your traceroute as a straight line or great circle in your animation. But it is actually something else.

And I'm wondering, have you all looked into actually finding some way outside? I mean, traceroute won't show you this, but finding some way to find out what actual geographic path is being used between the routers that show up in your traceroute?

I mean, the obvious question is, just because two routers that showed up in your traceroute are both in the same country doesn't necessarily mean that the path between them was the shortest path and didn't go all the way around the world or something, through many countries to get between the two routers. Have you looked into that at all?

>> ANTONIO GAMBA-BARI: Thank you for this question. I remember the first time we started with Google Earth and started looking at these various nice and smooth lines that become sort of misleading sometimes because as you pointed, we have tremendous challenges in terms of geolocation and geoprecision. Some of these questions come in terms of, where is it? At the building level? Is it the router at your house?

So we do have some question or issues about geoprecision. Some of the databases we start with are MaxMind. Some of you might be familiar. We rely on crowdsourcing, people telling us, "I know that router is there because it is my house," or some building levels like UFT, we run from that.

But despite precision of our geolocation, our geoprecision has been sufficient for the goals or the aim of our project which is, at this point, the city level will be enough for us to say that packets are going from Toronto via New York and then to Montreal.

So we don't have a solution for that. I know that Andrew and other of my colleagues, Colin over there, he's working in that particular aspect of geolocation.

But, as I said, we're welcome for comments or suggestions to how we can do that better. Thanks.

>> RUEDIGER VOLK: Ruediger Volk, Deutsche Telekom. For Antonio and the CAIDA research, in particular for the more developed parts of the

Internet, I wonder how much you're actually presenting the full picture.

Is it important to know whether a certain router is in this block on the Toronto map or on the other side of the town? Not quite sure. In Germany, I would know certain things that are closer to certain installations which might be something interesting. But actually does the IP routing trace really show how the traffic flows? How much traffic is actually directed by users who sign up with an email service somewhere across the globe? And is that stuff that is actually pretty much under user control -- they don't care unfortunately how much traffic -- how much traffic that is still routed in interesting ways is to be observed at the IP routing layer? And how much is actually happening more on the application side?

>> AMOGH DHAMDHERE: So I'm not sure I have a great answer to that. We do traceroutes. We don't really know how much traffic is flowing on those parts. Perhaps one way to look at those kind of questions is to do traceroutes to the kind of content you care about. You try to identify the CDN locations where traffic is coming to you from and look at the pass-through to those CDN locations.

And if you've managed to get the locations of the top content, then maybe you have a better answer to the kind of question you asked.

>> JANE COFFIN: Another idea to help with that -- and this is just a suggestion -- is that when you live in certain developing countries, you don't have an option for a very consistent service that you have a lot of confidence in. I know this from having lived in certain places. And you use those external services.

You do try and build local environment, and you host more services locally. You could take a look through what the project we're doing in Bolivia, take a look at the snapshot of the market now, where the traffic is being routed or where we know the traffic is flowing, the different traffic levels out of the country across the borders.

And then after the IX is installed, take a look at what those traffic level differences are, the quality of service levels, where latency and do some other corollary activity. We could actually take a snapshot there.

And I know some of the research that EuroIX and others are doing, they are trying to take a look at that as well. So that may be something where we start to take a look in those emerging markets to try to take a stronger perspective and do more analytics by putting in the servers, looking at the measurements, working with the community preIX, when the

IX is installed, and after so that we can help with that data correlation.

>> ANTONIO GAMBA-BARI: I would like to add just a quick note on the question about how much traffic goes on the application level or the user level.

For us, it's a very interesting question because currently our traceroute generation application is a crowdsourced application, so it's run on a specific port and it's run on demand by the user.

In the next level of our traceroute generation, we are thinking on more like a background application that could be monitoring the tremendous amount of services like happening when you are going to services like Gmail, Amazon.

So there is tons of exchange that's happening in the background of your computer. And having said that, part of the aim of our project is to raise awareness about that privacy or potential vulnerability of your privacy that is going on behind the scenes.

>> ANDREW SULLIVAN: So I'm conscious of the time, and because we have the IAB tomato-throwing event up next, I'm going to cut the lines now.

>>CARLOS MARTINEZ: Thank you. Carlos Martinez, LACNIC. This is a question/comment for Antonio.

I really appreciate your work and I think it's very interesting, very relevant. I do have an uneasy feeling about the use of the term "efficiency," because "efficiency" is something that is tied to a metric or indicator. You have to make things efficient in terms of delay, hidden costs in terms of red tape, for example.

So someone could argue with you that the current path traffic takes to get from BC to Toronto is perhaps the most efficient one, in terms of regulatory landscape, cost of laying fiber, things like that. Things that go beyond strict efficiency in terms of delay or distance.

So, my point here is that the perceived inefficiencies in how traffic is routed have underlying causes that go beyond technology, and those underlying causes should be addressed as well. Because if you change that boundary conditions, you can get the system to fix itself, in a way. Thanks.

>> JOEL JAEGGLI: Joel Jaeggli. I think my observation sort of dovetails with that, which is that there are geopolitical and economic

considerations here that created these pathways and they're pathways that are used for things other than Internet transit, for example, and so the implications of those sort of have long-lasting consequences.

I mean, if we ask why it is that the path from Chicago to Vancouver is about 900 route miles shorter, give or take, than the one from Toronto to Vancouver via Edmonton, you know, we don't really have to wonder about that.

But we can blame the Ashburton Treaty for it, which there are a number of Canadians who blame that for all sorts of things, right? Because the British sold you out.

But, yeah, I mean, there are some physical physics there that people who like rapid phone calls and computer games that run fast are really in favor of that are actually sort of hard to defeat without redrawing lines or not worrying about them as much.

So I don't really have to think very hard when I wonder why traffic goes to Chicago or New York from places that have very dense economic ties with those locations, right? Trucks do. So do ships and people in cars and airplanes and so forth. So it's not just the packets.

>> UNIDENTIFIED IETFer 1: Jane, one of the more interesting metrics I thought you presented was the cost per unit balance, and I'm just wondering: Has any effort been done to sort of map this to geographies and to use that as a way to sort of focus attention on particular problem areas?

>> JANE COFFIN: My colleague, Michael Kende, who is an economist for the Internet Society, has done some mapping and some corollary work. We're going to be working more together to do some of that, as you're suggesting.

We've got some mapping together now of the IXs, but we also want to take a look at the corollary cost considerations where costs have come down, where latency has come down, and analyze it in greater detail to look at some of the other factors per region, per continent, and what we can divine from that. Because obviously, some of this is regulatory and policy, some of this has to do with a strong technical community, and look at those different factors.

It can help influence how you go into a certain region, how you work with a certain community, and try to push-pull on different factors.

We could be more efficient, actually, if we did some of that, and I

think that's definitely in the near future. It's just putting all of that together with an economist, with the technical experts, and getting some of that data.

It is hard to get the data. I'll be honest. And that's part of the objective of what we're doing in this toolkit, but also all the colleagues around the world are trying to do the same.

I will say the RIPE ATLAS project has some great data. So if we can start to think of ways to correlate some of what Amogh may be doing and you, Antonio, and the others to bring some of those data sets together, we could find some better answers to that question and try and be more efficient with what we're doing and how we do it.

I will say one thing about speed. I'm reading a book called "Flash Boys" that was recommended to me by Phil Roberts, who I know is in the room.

If you think about, in the highly developed markets, what they're trying to bring down the speeds by -- I said, "Oh, seconds?" He said, "No, Jane, we're looking at milliseconds here." And I said, "Right, right."

But if you think about the new network topologies where they're trying to put in a straight line between Chicago and New York in order to trade faster, that's a whole different issue than what we're trying to do in Kinshasa, where, quite frankly, after the IX was up and running, the mobile operators found in four hours that the quality of service had gone up for their network delivery, for their traffic delivery.

So you've got humongous different aspects here, with developing countries and more developed, but we can learn something from each of those experiences for sure.

>> ANDREW SULLIVAN: It does strike me, just picking up on what Joel said before, that it would be interesting to take some of the data that you've got now -- for instance, the CAIDA data, the IXmaps data -- look at that and then track the new IXs going in and say, "Well, is that actually changing the way these patterns are?"

Because, Joel's point that, well, people are following the patterns that they ever did. If there's a new road there, do people change direction? Because they don't always, right? I mean, just in physical space. Anyway...

>> UNIDENTIFIED IETFer 2: So the first one is a kind of a comment and a

question following that.

I was wondering how much of this strange routing is due to the desired aggregate, the point where the carriers have so-called free arrangement of exchanging traffic on both directions. Maybe those are the points they want to go and they won't aggregate enough traffic so that they can utilize that whole thing.

And my question kind of dovetails to that, that if you have two carriers who don't have the same amount of traffic, let's say, you know, probably like everyone knows that Netflix uses Level 3 as one of their providers, and that peers with Verizon, and then of course, you know, Netflix is pushing an enormous amount of data compared to what will go upstream from there and back to Level 3. So do you have this kind of problem where carriers don't want to peer because they would not get a very good economic arrangement?

>> ANTONIO GAMBA-BARI: I'm not sure we caught all your question. Could you rephrase quickly your question, very short?

>> UNIDENTIFIED IETFer 2: Well, I'm basically saying some of the peering points have economic advantage because let's say Carrier A is talking to Carrier B at Point B, which could be in Chicago, for example, but they have an arrangement that they can freely exchange data, so they will try to move all the data through that, regardless of the actual distance between the two final destinations, right?

And if you cannot create such a partnership with the two carriers, will they then be encouraged or will they be even willing to create a peering relationship amongst themselves, when they are trying to encourage more IXPs?

>> JANE COFFIN: Some of the more advanced networks run by an incumbent don't necessarily peer at the Internet exchange point. In Johannesburg, which is a 16-year-old IXP, the incumbent isn't peering at the IX, so there are some smaller networks which may not be as sophisticated but they level-up, if you will, from a technical perspective as there is more confidence in their own infrastructures and making more money and they're exchanging traffic. It's just a network effect from that peering at the IX.

Incumbents aren't necessarily necessary, and I guess the point you're making is that you have to have two really great networks in order to send your traffic.

Well, that may be true if we're running stock market traffic, but if

you're just talking about just getting better quality of service, lower latency, more efficiency from the network engineers who are learning how to better route traffic, quite frankly, some aren't deploying BGP and they need to understand how to do some of that engineering, and we try and help there. But those network efficiencies do come about after some time. And that may be where those sophisticated networks grow and become better in what they're doing.

>> ANDREW SULLIVAN: So I did close the line but there if you can make it short, you can go.

>>ROBERT KISTELEKI: I can, yeah. Robert Kisteleki, RIPE NCC.

You said wouldn't it be nice to start collecting data before, during, and after the IXP deployment.

We are exactly doing that, together with ISOC. So we started that thinking process at the minimum. We have deployment. We also have deployment in Africa, which is difficult to do, but we could and did start local measurements, you know, whether traffic is within the country or not, and we would like to help ISOC figuring it out where are the best spots for IXPs and how the things change if you actually have an IXP locally. So we are on that, together with ISOC.

>> ANDREW SULLIVAN: Excellent. Thank you.

All right. I want to thank Antonio, Jane, and Amogh. Thank you very much for your time today, and with that, I'm bringing the rest of the IAB up. Yes. Thank you very much.

[Applause]

>> RUSS HOUSLEY: Okay. So Andrew is over there. He is the one that asked for the tomatoes.

We'll start down here with Mary for the introductions.

>> MARY BARNES: Mary Barnes, IAB.

>> JARI ARKKO: Jari Arkko, IETF chair.

>> TED HARDIE: Ted Hardie, IAB.

>> LARS EGGERT: Lars Eggert, IRTF chair.

>> XING LI: Xing Li, IAB.

>> ELIOT LEAR: Eliot Lear, IAB.

>> DAVE THALER: Dave Thaler, IAB.

>> RUSS HOUSLEY: Russ Housley, IAB.

>> ERIK NORDMARK: Erik Nordmark, IAB.

>> JOE HILDEBRAND: Joe Hildebrand.

>> MARC BLANCHET: Marc Blanchet, IAB.

>> JOEL HALPERN: Joel Halpern, IAB.

>> BRIAN TRAMMELL: Brian Trammell, IAB.

>> ANDREW SULLIVAN: Andrew Sullivan. You already know me.

>> HEATHER FLANAGAN: Heather Flanagan, not IAB.

>> RUSS HOUSLEY: Okay. The mics are open.

[No one at the microphone lines.]

>> RUSS HOUSLEY: Going once? Going twice? Enjoy your dinner.

[Applause]