

Addressing Global Routing as a Community

Brenna Meade
Network Engineer
International Networks at Indiana
University



**International
Networks**
at Indiana University

What is BGP

- BGP or Border Gateway Protocol is the protocol used between routers to exchange routing information and reachability information between or inside an AS on the Internet.
- BGP makes the Internet work, and in most cases it just works
 - Needs to be tuned for best performance
- BGP makes routing decisions based on paths, network policies and rule-sets, etc.

BGP JUST WORKSuntil it doesn't

R&E Routing Architecture Vs. Commodity.

- Research and Education Networks
 - Bandwidth
 - Performance Engineering
 - Deterministic behavior
 - Community
- Commodity Networks
 - Traffic shaping
 - DoS protections
 - Unknown architecture
- R&E networks are engineered to support science while commodity networks are not
 - Keep the science traffic on the science networks!

Takeaways!

Routing will not take care of itself

- Old routes may not work well with new networks
- New routes may not work as planned

How do we address routing anomalies as a community?

The Routing Working Group!

How we used to solve this:

I think I've found a bad route. How do I address this?

Start a conversation! "Hello, it appears that there's a lot of traffic going between your institution and X that probably shouldn't

Can we work together to improve it?"

International Routing Working Group

Established to create a community to work on international routing issues to improve R&E performance. The goal is to engage network owners and NRENs to not only reactively discuss and address ineffective routes, but will work proactively across the community to systematically create policies to prevent them from occurring.

- Co Chairs:
 - Brenna Meade , Indiana University
 - Warrick Mitchell, AARNET
 - Hans Addleman, Indiana University

What do we do now? - The Routing Working Group

Routing Working Group case process

1. Cases are submitted via the mailing list, slack or at the monthly meeting
2. Teams are selected to assist with the case by the chairs
3. The case is added to the master case list (open access)
 - [RWG Master Case List](#)
4. A folder is created for the specific case and team members are given access to the documentation
 - Most cases are worked on via email or slack
5. Case updates are given at each RWG meeting

Routing Working Group - What are the goals?

- **Engineering focus**

- Document possible erroneous routes
- Identify teams to address them
- Check in together as we work through them

- **Policy Focus**

- Detail routing policies for paths
 - Including preferred backup paths!
- Verify if policy is being followed

Routing Challenges We've Observed

Asymmetrical routing - meaning a source to a destination takes one path and takes a different path when it returns to the source

R&E data takes a less efficient route around the world - affecting performance

- Europe to Asia routes traversing the US
- Africa to Europe routes traversing the US

R&E data takes a commodity route when an R&E path is available

New R&E links are removed or added but routing does not adjust appropriately

Leaking of Private ASN's into the global routing table by R&E networks

IP blocks advertised with a Bogon Origin ASN's within R&E routing table

Tools Available

- NetSage gives quick access to Flow, SNMP, and route information
 - <http://portal.netsage.global>
 - Easy to find data on rate and volume
 - Clues to routing information as well
- Traceroute
- Routeviews peering
- Router Proxy / Looking Glass

ITB <> Starlight

Summary: Asymmetrical routing is preventing access the starlight ESnet DTN.

- Route from ITB to ESnet Chicago is using R&E networks via TEIN, TransPAC and Pacific Wave
- Chicago ESnet to ITB is using the PCCW and commercial paths.

Resolution: Worked with ESnet to update accepted routes from TransPAC to use R&E path. Trace routes are symmetrical now, We received confirmation from ITB engineers that they are able to access the DTN.

Team: Simon Peter Green (SingAREN) , Basuki Suhardiman (ITB) , Kate Robinson (ESnet)
Brenna Meade (IU)

Singapore to New Zealand

Summary: Traffic from Singapore to New Zealand traversing the US, and is asymmetrical

Resolution: Asymmetric routing has been resolved by moving traffic to the Singapore <> Guam link and confirmed via perfSONAR tests for both IPv4 and IPv6

Singapore to New Zealand

Routes before changing the routing:

SingAREN@SG -> SingAREN@LA -> Internet2@LA -> REANNZ

SingAREN@SG -> APAN-JP -> Transpac@SEA -> REANNZ

Current symmetrical routing (using Guam <> Singapore link) :

SingAREN@SG -> GOREX-> REANNZ

Team: Brenna Meade (IU) , Dylan Hall (REANNZ) , Francis Lee (SingAREN), Simon Green (SingAREN)

Singapore <> Taiwan

- Route from Taiwan to Singapore traversing the US
- Team: Brenna Meade (IU) , Simon Green (SingAREN), Eric Yen (Taiwan), Wenshui Chen (Taiwan)

Before - asymmetrical routing

Taiwan to Singapore

1. ASGCNET (Academia Sinica Grid Computing Center, Taiwan)
2. ASGCNET - INTERNET2 exchange , Chicago
3. Internet2 Chicago
4. Internet2 Kansas
5. Internet2 Salt Lake
6. Internet2 LA
7. Internet2 LA - SingAREN LA exchange
8. SingAREN LA
9. SingAREN
10. SingAREN Final

Taiwan to Singapore takes route through U.S. via Chicago / TWAREN, and SINGAREN LA

Singapore to Taiwan

1. SingAREN
2. SingAREN - APAN JP exchange
3. APAN JP - ASGCNET exchange
4. ASGCNET

Singapore to Taiwan does not take U.S. route

After: Singapore to Taiwan route is no longer asymmetrical

Taiwan to Singapore

1. ASGCNET (Taiwan)
2. ASGCNET - APAN JP
(Tokyo) exchange
3. APAN JP - SINGAREN
(Singapore) exchange
4. SINGAREN

Singapore to Taiwan

1. SINGAREN
2. SINGAREN - APAN-JP
exchange
3. APAN-JP - ASGC
exchange
4. ASGCNET - TAIWAN

Internal to Asia traffic traversing the US

Problem: seeing Hong Kong to China traffic in US

Netsage: Beijing to CUHK over transPAC

Status: BGP peering was adjusted between CSTNET and HARNET, traffic is no longer taking transPAC

Team: Brenna Meade (IU) , KW Pong (ITSC) , Doug Southworth (IU)

Tool Talks

MANRS - [MANRS - Presentation Slides](#)

Router Proxy / Looking Glass - [Router Proxy - Presentation Slides](#)

Peering DB

- [Peering DB - Presentation Slides](#)
- [PeeringDB Video](#)

Tool Talks

Perfsonar

- [Perfsonar - Presentation Slides](#)
- [Perfsonar - Video](#)

MicroDep - [MicroDep - Video](#)

Netsage - [Netsage and Finding Weird Routes - Presentation Slides](#)

Developing a Collaborative BGP Routing / Analyzing / Diagnosing Platform -
[Developing a Collaborative BGP Routing Analyzing and Diagnosing Platform](#)

Current Status

167 members are subscribed the RWG list

- Over 20 countries represented
- Over 80 institutions represented

14 total cases

10 solved cases

2 cases on hold due to lack of resources or pending new links

Goals for 2023

Continuing working on routing cases with the community

Provide more tool talks

Create a series of community driven BGP best practice documents

- Best practices for R&E network providers specifically
 - how do they differ from other networks?
- Scalability
- Load balancing
- R&E Peering Agreement best practices
- BGP Communities - how to use this with a focus on R&E networking

Upcoming Meetings

Meetings are the 3rd Tuesday of every month

Upcoming Meetings:

January 17th, 2023 - 2023 Kickoff

February 21st, 2023

March 14th, 2023 - In Person / Remote APAN

If you have suggestions for tools talks or an interest in presenting at the Routing Working Group, let one of the chairs know!

Submit your cases!

Email the Chairs!

meadeb@iu.edu

addlema@iu.edu

warrick.mitchell@aarnet.edu.au

Join the routing working group!

Mailing list routing-wg@gna-g.net

- Contact Brenna to be added meadeb@iu.edu

Slack

- APAN Slack Instance, Channel: Routing

Web

- <https://www.gna-g.net/join-working-group/gna-g-routing-wg/>

Contact any of the co-chairs for more information!