

## Follow-up to Observations on Anycast Topology and Performance

Steve Gibbard Packet Clearing House



## **Original paper**

- Observations on Anycast Topology and Performance
  - Presented at last OARC/DNS Ops meeting
- Looked at query distribution from several anycast systems
  - > J Root data from Verisign
  - K Root data from RIPE
  - > C, F, K Root data from CAIDA



## Original paper (cont.)

- Observed server selection following customer relationships rather than geography for J and K roots
  - Transit from different providers at different nodes
  - ISPs prefer to send traffic to customers, regardless of geography



# Original paper (cont.)

- Attempted to show that an anycast system being consistent about transit and peering policies avoided those issues
  - Did data collection from PCH anycast system
  - > Four global nodes, transit from NTT/Teleglobe
  - Some peering on the global nodes
  - Traffic followed geography, except when it didn't
    - Cases where it didn't were anomalous



## **Original paper -- questions**

- Methodology issues
- > Anomalies in the data
- Local node performance



## Methodology issues

- Used unique query sources (actually source /24s) instead of hit counts.
  - Did that distort results, or miss important networks?
- > Reran analysis, using hit counts, for answers:
  - > Results mostly similar.
  - Weighting by query sources did have interesting effects. Half of "Belarus" queries from one network in the US.



### Local nodes

> Left out of earlier analysis:

- > Initially considered out of scope.
- > Added in later tests.
- > Results more or less as expected:
  - Isolated regions -- Nepal, Bangladesh, Kenya, mostly self-contained.
  - More major nodes -- Amsterdam, Stockholm, Singapore, drew from wider, but geographically contiguous, areas.



### **Data Anomalies**

- Traffic anomalies, blamed on inconsistent peering:
  - Southeast Europe/North Africa to Ashburn
  - Indian sources to Ashburn
  - > Spanish sources to Hong Kong
  - Lots of Asian traffic to Palo Alto rather than Hong Kong



## **Chasing anomalies**

Southeast Europe/North Africa
210k queries from Telecom Italia
62k queries from UPC
We peered with both in US but not Europe.
Turned up peering, and traffic shifted.



### **Peering shift**



#### Ashburn, before shift



#### Ashburn, after shift



#### Amsterdam, before shift



Amsterdam, after shift



### Asia to Palo Alto

- Much of Asia hotpotatoing to Palo Alto in January
- > Still doing so in March
- Going to Hong Kong in May
- Spot check shows these mostly seen through transit, so shift presumably in transit provider networks



Palo Alto, January



Hong Kong, September



### New issues

Turk Telecom --- Uses circuit to Eastern US. Not much we can do in short term.
Planned node in DE-CIX might fix this.
Things that changed on their own:

No longer seeing Jazz Telecom in Hong Kong -- not sure why.
Indian traffic heading East rather than West.



### Latin America

- Going to right global node. Local nodes are closer.
- Could be fixed by peering in Miami or Sao Paulo
- > Ending up in Ashburn:
  - Region-wide: 474k qpd from Telefonica.
  - Caribbean: 142k qpd from Columbus/New World
  - Brazil:
    - > 150k qpd from Bahia
    - > 80k qpd from Telefonica
    - > 80k qpd from Embratel
    - > 60k qpd from Brazil Telecom.





## Conclusions

- Fixing peering issues does fix performance issues.
  - Some networks cooperate with this, while others make it difficult.
- > A global node closer to South America might be useful.



### Thanks!

### Paper at:

http://www.pch.net/resources/papers/anycast-performance/

Steve Gibbard Packet Clearing House www.pch.net scg@pch.net