

# An analysis of the Internet interconnection density in IPv6 compared to IPv4

Christian Kaufmann  
@DENO5 in Darmstadt

# Master Thesis- Finding a Topic

- Studying again
  - > Master of Science in Advanced Networking
- Daily life problems, relevant to a lot of people
- Inspiration from Martin Levy
- An analysis of the Internet interconnection density in IPv6 compared to IPv4

# So what is the issue?

- IPv6 latency / speed is not the same as on IPv4
- Because of:
  - MTU, tunnels, hardware, etc.
  - The peering interconnection density in v6
- The amount of interconnection density seems different in v4 and V6 for various reasons:
  - V6 is still in a test phase
  - Just new peering sessions are dual stacked
  - Etc...

# How to measure and proof it...

- Network latency is measured with pings and traceroutes
- Lets take a lot of them from a lot of sources in IPv4 and IPv6 to common destinations
  - RIPE ATLAS as the weapon of choice
  - 500 sources to 500 destinations
  - For v4 and v6
  - Use traceroutes instead of ping to get more info
  - Measuring RTT, IP Hops and ASN Hops
  - $500 \times 500 \times 2 \times 3 = 1.500.000$  Data Points

# Sources - Sampling

- RIPE ATLAS Probes - <http://atlas.ripe.net>
  - Over 4000 globally distributed probes, API, service for free
  - 500 max amount probes against one target at a time
  - Scope: global
  - Dualstacked: A and AAAA record
  - Probe ID feature secures consistency of the sample to ensure comparability of the measurements

# Destinations – Sampling

- Websites – to make our life easier
- Source: <http://www.alexa.com>
- Requirements for sampling:
  - Dualstacked
  - Nicely globally distributed
- Handpicked 500 (to have a nice 500 by 500 matrix😊)
- NO CDN's or anycast deployments – in order to measure the interconnection density not the website performance

# Number crunching – Collecting Data

- Script 1:
  - Triggers the test ‘500 fix probes against 1 destination’
  - Runs 500 times
- Script 2:
  - Download 1000 JSON files with the 500 traceroutes
- Script 3:
  - Takes the last hop RTT of the traceroute
  - Records the IP Hops of every test
  - AS Lookup for every IP Hop (takes for ever 😊)
  - Create a CVS / Excel output

# Issues on the way...

- 😊 Huge Sample, Flood of Data
- RIPE Atlas does not like too many tests at once
- Not all probes are online all the time
- Not all websites are online all the time
- ICMP filtering seems to be more common on V6 then V4 - especially on Hosts
  - suddenly no data to report on
  - Little AS work around and some assumptions



# Next.... – Analysis and Interpretation

- Crunch this 1,5 mio data sets and throw the outliers and broken ones away
  - Analyze data– come to some kind of conclusion
  - Hypothesis: A lower number of IPv6 than IPv4 peering sessions exist as of today
  - If the hypothesis will be confirmed by the data analysis: Ask the networks for the reasons for this imbalance
  - Outline a solution which will help to push the number of IPv6 peering sessions
- > Write a Thesis 😊
- Report back to the community

# Thanks to:

- Martin Levy for the idea and inspiration
- Vesna, Kaveh and Robert from the RIPE NCC
- Steven Schechter for the “little” programming 😊

# Questions and Comments?

- Is this research actually helpful for the community?
- Does anyone else experience the latency issue and lack of peering as an issue?
- Email: [ckatminxsdotnet](mailto:ckatminxsdotnet)