Round Trip Times Between Resolvers and a Root Server

Paul Hoffman

DNS-OARC
July 2022

v4

Overview

- Resolver operators desire for better round trip times (RTTs) for their root service
- Useful to compare proposed RTTs to those seen by other resolvers
- Estimating RTTs to IMRS instances by measuring IP time to live (TTL) and comparing it to RIPE Atlas measurements
- Request for other RSOs to make similar RTT measurements

The desire for faster root service

- **RSOs are often asked** to place anycast instances in particular locations in order to bring faster root service to a group of resolver operators
- “We want our resolvers to have at least typical performance for root service”
  - **Median** of resolvers across the RSS
- “We think our resolvers get about the worst performance for root service”
  - **90th percentile** of resolvers across the RSS

Estimating RTTs to IMRS instances (1)

- An easy design would have been “sample the queries we get, ping a sample of the querying addresses, measure the RTT”
- From the IMRS FAQ: “ICANN just collects data and telemetry from the server about how many and which queries it gets. **No other data is available to or collected by ICANN**.”
- But that doesn’t stop us from estimating!

Estimating RTTs to IMRS instances (2)

- Like most (all?) RSOs, IMRS collects its query traffic
- This traffic contains the DNS queries, but it also contains the address of the querier and other data from the IP header, including the **IP TTL value**
- The IP TTL can be an approximate proxy for the RTT
- Take a full day's worth of IMRS traffic for all instances
- Mash and extract, focusing on the IP TTLs

IMRS TTL results

- A total of 1,737,156 unique **obscured** addresses were seen across all the instances on that day
- **Median IP TTL is about 10.5 hops** for both raw and weighted calculations
- **90th percentile IP TTL is about 18 hops**
- Note that forward routes and reverse routes might be different and have different IP TTLs

Using RIPE Atlas to correlate IP TTLs to RTT

- Use RIPE Atlas to ping the IMRS anycast v4 and v6 addresses from about 10,000 probes
- Map the RTTs to the IP TTLs from the probes
- A median IP TTL of **10.5 hops is about 33 milliseconds**
- A 90th percentile IP TTL is about **18 hops is about 60 milliseconds**
- Notes
  - Probes are not at the same place in the network as typical DNS resolvers, so this is a bit of a leap of faith
  - The Atlas numbers get more chaotic around 18 hops

Is this “good”?

- This set of measurements doesn’t say whether the median or 90th percentile values are “good enough” for typical DNS root service
- Faster is better, but can it be perceived by the users of the resolver?

A request to other RSOs

- **Please find the median and 90th RTT percentiles** so we can compare latency of RSOs across the RSS
- You can hopefully measure directly from sampling of your traffic, but if not, you can probably do something like what we did here
- **We can share results**, maybe have a joint publication