

Comparing Timeout Behavior and Reachability at I-root and L-root

Introduction and Background

This project investigates timeout events observed in RIPE Atlas DNS measurements toward I-root, operated by Netnod, and L-root, operated by ICANN. The goal is to determine whether these events are mainly caused by probe-side issues, anycast catchment shifts, or path- and root-specific reachability problems.

The same root may be reached through different physical instances depending on routing. This makes timeout interpretation difficult, since a timeout does not directly reveal where the problem occurred. We therefore combine cross-root comparison, probe metadata, anycast instance names, and traceroute data to identify which probes are affected and how timeout events should be interpreted.

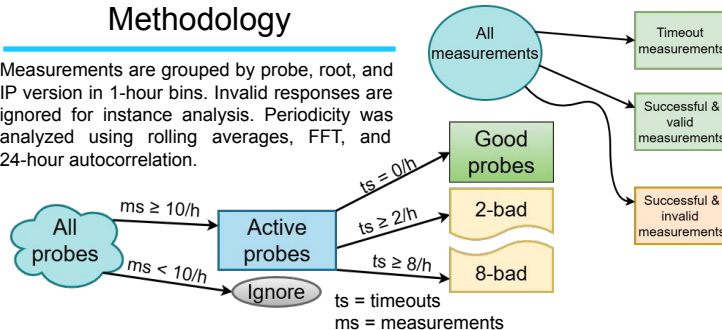
Research Questions

How do timeout behavior and reachability differ between DNS root servers I and L?

- 1) Do timeout rates toward I-root and L-root exhibit periodic behavior, and do these differ between IPv4 and IPv6?
- 2) Which probes experience timeouts toward both roots, toward only one root or are always failing?
- 3) Which root on which IP version is more stable from the perspective of the number of probes that timeout?
- 4) Are the timeouts related to client-side factors, server-side behavior, or network paths?
- 5) How do sudden changes in observed root-instance catchments affect DNS reachability, and can probes be redistributed without visible timeout loss?

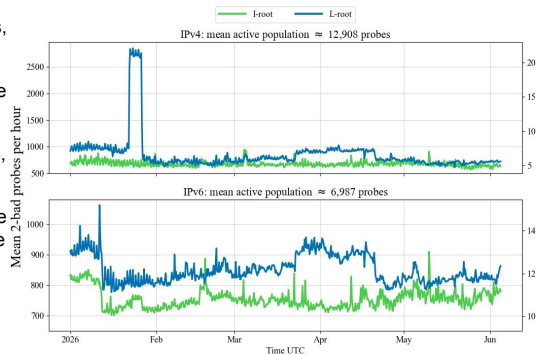
Methodology

Measurements are grouped by probe, root, and IP version in 1-hour bins. Invalid responses are ignored for instance analysis. Periodicity was analyzed using rolling averages, FFT, and 24-hour autocorrelation.



Main Results

- L-root showed stronger timeout variation than I-root, especially over IPv4;
- IPv6 failures were more often shared across both roots, while IPv4 showed more root-specific timeout behavior;
- During the L-root IPv4 event, many traceroutes stopped right before reaching L-root, suggesting reachability problems near ICANN's peering edges.





Limitations

- IPv4 and IPv6 probe population differ; IPv6 has much fewer probes involved in measurements;
- RIPE Atlas global view is large, but probes are unevenly distributed across countries and networks;
- DNS timeouts are ambiguous as they do not reveal the failed component or the anycast instance that would have answered;
- Traceroutes and probe metadata help with interpretation, but both can contain missing, incomplete, or stale information.

IP version	Probe-hour state	Mean probes/hour	Mean share (%)
IPv4	2-Bad toward both roots	306	2.37
IPv4	2-Bad only toward I-root	356	2.76
IPv4	2-Bad only toward L-root	524	4.06
IPv4	Fully failing toward both	248	1.93
IPv6	2-Bad toward both roots	670	9.59
IPv6	2-Bad only toward I-root	76	1.09
IPv6	2-Bad only toward L-root	169	2.41
IPv6	Fully failing toward both	609	8.71

Series	All instances seen	Median/day
I IPv4	96	91
I IPv6	86	80
L IPv4	134	108
L IPv6	99	82



Conclusions

- Failures toward both roots suggest shared probe/path issues, while failures toward only one root suggest root- or path-specific reachability problems.
- RIPE Atlas data can reveal significant reachability events and narrow down their causes, but stronger attribution requires analyzing external sources.
- Low to medium periodic timeout patterns were observed, but they were usually distributed across many probes rather than caused by one fixed failing group. Already unstable probes became worse during periodic episodes.