Round-Trip Times Between Resolvers and IMRS

ICANN Office of the Chief Technology Officer

Paul Hoffman
OCTO-036
7 June 2023

ICANN

TABLE OF CONTENTS

This document is part of ICANN's Office of the Chief Technical Officer (OCTO) document series. Please see the OCTO publication page for a list of documents in the series. If you have questions or suggestions on any of these documents, please send them to octo@icann.org.

This document supports ICANN's strategic goal to strengthen DNS root server operations governance in coordination with the DNS root server operators. It is part of ICANN's strategic objective to strengthen the security of the domain name system and the DNS root server system.

Executive Summary

The “IMRS Instance Placement Study,” OCTO-018, showed how ICANN can measure the number of Internet Protocol (IP) hops between resolvers and ICANN Managed Root Server (IMRS) instances. The goal of the research described in this document is to inform the community about how long it takes for resolvers that are “average” and “far” from an IMRS instance to get answers to their queries. This study focuses on measuring expected performance for resolvers that connect to IMRS.

1 Introduction

Domain Name System (DNS) resolver operators often want better round-trip times (RTTs) to the root server system (RSS) for root service. Basic access to the RSS does not cost them any more than normal traffic, but they may want to pay to put their resolver close to a root server operator’s instance. It is easy to see that some resolvers get faster responses to queries due to their location in the Internet’s topology.

Because of these different RTTs, it is useful to compare the RTTs seen by one resolver to those seen by others. Resolvers tend to use one or more root server operator (RSO) that gives a fast RTT, so it is useful to make this comparison of RTTs per RSO. Note that some resolvers do not always pick the fastest RSO: many use software that picks one of the fastest if they have similar speeds. A resolver operator, of course, would want all the RSOs to have better RTTs for their resolvers to always get faster service.

The IMRS is one of the 13 sets of servers that provide authoritative answers to queries in the global DNS. Like all RSOs, IMRS uses *anycast routing* to make many individual servers look like just one server with an IPv4 and an IPv6 IP address. In fact, IMRS has nearly 200 instances across the globe (as of May 2023), each of which replies to queries sent to either its IPv4 or IPv6 addresses. The main purpose of anycast routing is to cause hosts on the Internet (such as resolvers) to automatically connect with the closest instance without any configuration on the part of either the host or the instances.

This research is a follow-on to the research described in OCTO-018. While that study focused on finding locations that were furthest away from any IMRS instance, this study focuses on measuring expected performance for resolvers which connect to IMRS.

The research in this document extends that work by determining the correlation between the number of IP hops (the *hop count*) and the number of milliseconds it takes for a resolver to get a response from IMRS instances. The RTTs to IMRS instances are estimated by measuring the hop count, which is measured with the IP time-to-live (TTL) field in packets, and comparing it to RIPE Atlas measurements for systems with similar hop counts to IMRS instances.

Note that other RSOs can also measure this data for the same purposes. Because each RSO deploys instances differently, their measurements might come out differently than those presented here due to different routing configurations and different styles of network connectivity.

2 Data Collection Process

An easy design for estimating RTTs between resolvers and IMRS would have been "sample the queries that IMRS gets, ping a sample of the querying addresses, and measure the RTT of those pings." Formally, a "ping" is an Internet Control Message Protocol (ICMP) echo request. However, that is not permitted because it would violate the terms of IMRS's agreement with the operators of the IMRS instances. On the webpage, Frequently Asked Questions Hosting IMRS, it reads "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."

That restriction, however, does not stop us from using other measurements that can help model RTT measurements. IMRS logs the query traffic at its instances, then collects those logs on a daily basis. These logs contain the DNS queries and the address of the querier. Another piece of data in the logs is the IP TTL value from the IP header; this value is used in IP routers to measure the "hop count," which is the same as the number of routers that the IP packet has traversed. We believe that the hop count for the query going to the IMRS instance can be an approximate proxy for the RTT, if it can be combined with other data.

When collating the hop count data, the IP addresses of the querying resolvers are obscured. IPv4 addresses have their last byte turned to 0, and IPv6 addresses have their last ten bytes turned to 0; in technical terms, only the /24 (IPv4) and /48 (IPv6) is kept. For example, all resolvers with IPv4 addresses starting with "100.101.102" would be treated as a single resolver. The names in the DNS queries are discarded before collating the hop count data; this makes the dataset even more privacy preserving. The query names have no relevance to the study, so discarding them does not reduce the accuracy.

3 Data Analysis

The Hop Count Dataset

This data is called the "hop count dataset" in the analysis below. For this analysis, the hop count dataset is based on seven days of data from all the IMRS instances. The hop count dataset was collected in April 2023.

In the hop count dataset, a total of 3,228,757 obscured resolver addresses (both IPv4 and IPv6) were seen across all the IMRS instances in that week. As a side note, 9,085 (0.3 percent) of the addresses were in IPv4 or IPv6 private address spaces, and thus were dropped from the analysis.

Figure 1 shows the distribution of the number of hops between resolvers and their closest IMRS instance. The distribution is quite uneven, possibly due to the method that common resolvers use to pick which root server to use.

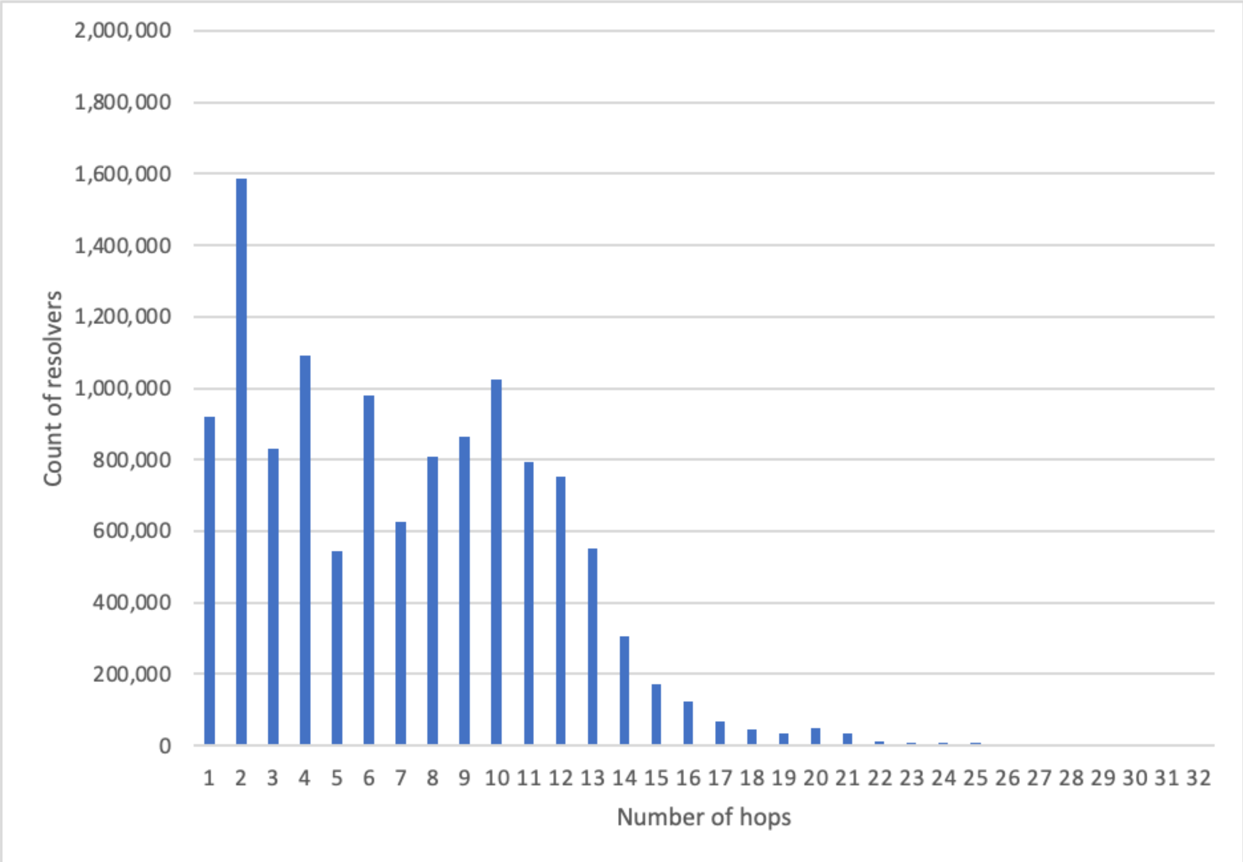

Figure 1. Number of hops versus the count of resolvers

For this analysis, the median and 90th percentiles of the data in Figure 1 are the interesting values. Figure 2 shows the same data, but measured cumulatively.

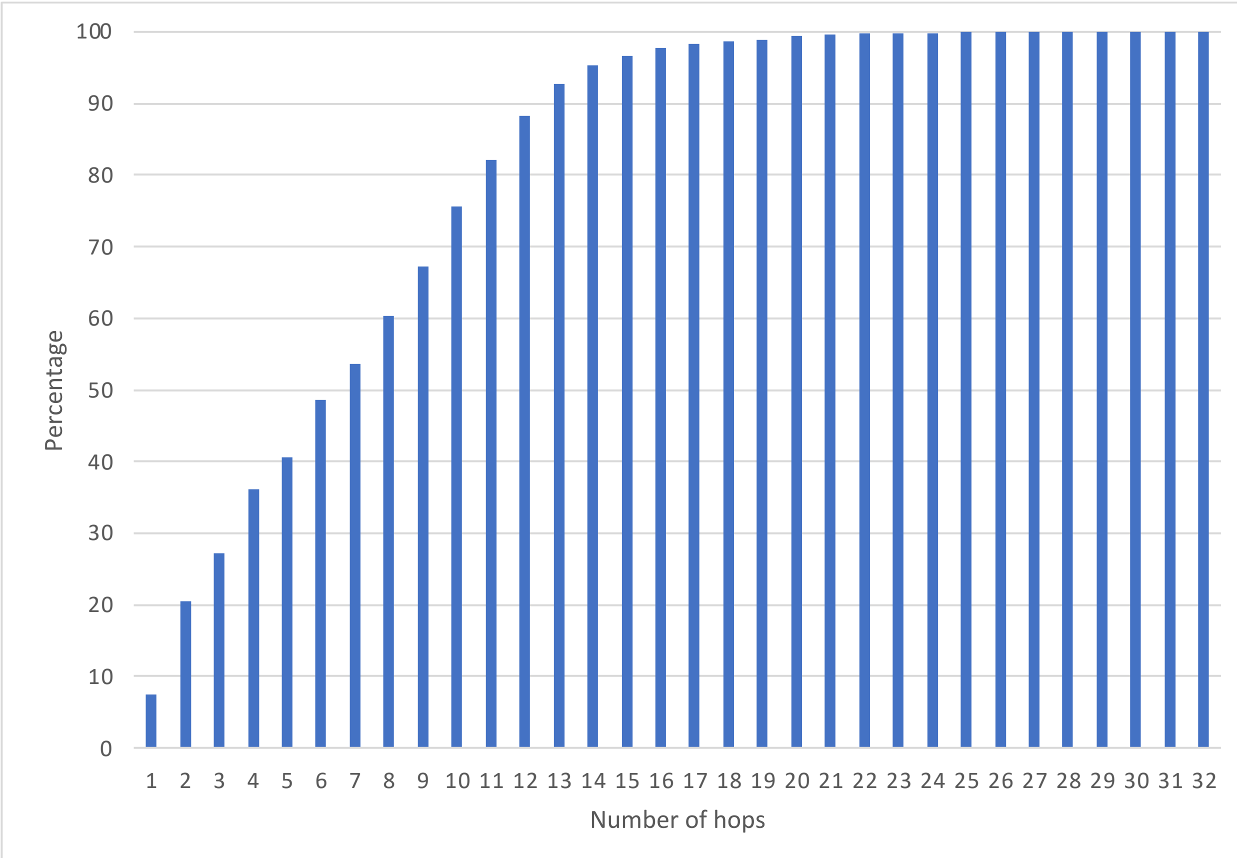

Figure 2. Number of hops versus the cumulative percentage of resolvers

The analysis of the data shows that the median IP TTL is about 6 hops, and that the 90th percentile IP TTL is about 13 hops.

3.1 The RIPE Atlas Dataset

In parallel with getting the hop count dataset, about 10,000 RIPE Atlas probes and anchors were instructed to send queries to the IMRS IPv4 address and IPv6 address. For each probe or anchor, both the hop count to the IMRS instance and the RTT of a simple DNS query were recorded. This data is called the “RIPE Atlas dataset” and was collected in April 2023.

Figure 3 compares the hop count to RTTs in queries sent to IMRS instances, as seen by RIPE Atlas probes.

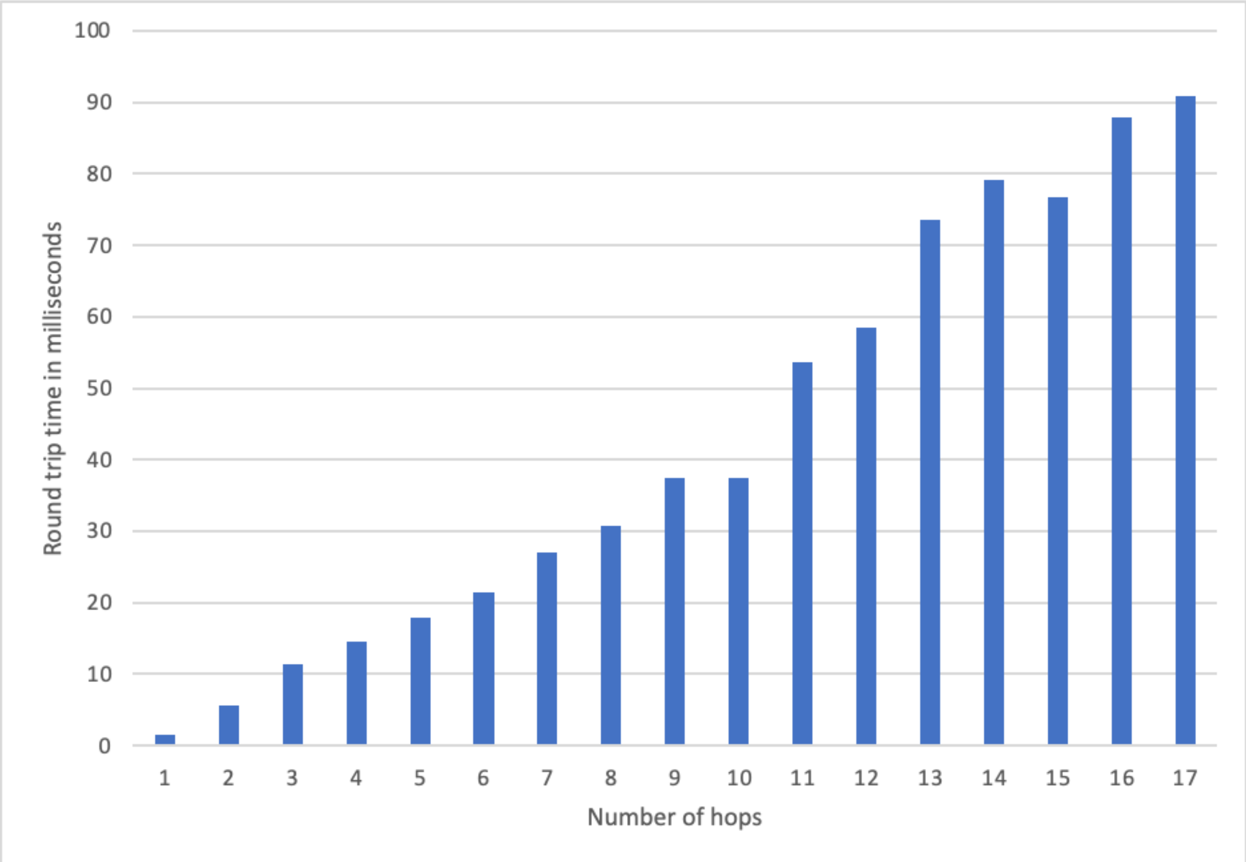

Figure 3. Number of hops versus round trip times

In making the analysis, note that the RIPE Atlas probes are not at the same place in the network as typical DNS resolvers, so using these numbers can lead to inaccuracy – if the RIPE Atlas probes are at significantly better-connected, or significantly worse-connected, locations than typical DNS resolvers. Also, note that RTTs get somewhat chaotic above 14 hops, which also makes the analysis less precise.

3.2 Combining the Datasets

A median IP TTL of 6 hops gives an RTT of about 25 milliseconds, while the 90th percentile IP TTL of 13 hops is about 70 milliseconds. This means that a resolver that has Internet connectivity typical for resolvers that use IMRS will have RTTs of about 25 milliseconds, and a resolver that has bad connectivity among resolvers that use IMRS will have RTTs of about 70 milliseconds.

4 Conclusion

We estimate that most resolvers typically get IMRS responses back in about 25 milliseconds, and resolvers with poor Internet connectivity receive responses back in about 70 milliseconds. These times are likely to be fast enough to be undetectable to users of those resolvers. As resolver operators think about how they should spend their resources to improve users' experience, increased speed to root servers could likely be quite low on their list of priorities. As seen in this paper, RTT performance is only one consideration for resolver operators.