

# DNS Root Server Equivalence Study

ICANN Office of the Chief Technology Officer

Roy Arends  
OCTO-042  
14 January 2026



---

## TABLE OF CONTENTS

<b>EXECUTIVE SUMMARY</b>	<b>3</b>
<b>1 INTRODUCTION</b>	<b>3</b>
<b>2 METHODOLOGY</b>	<b>3</b>
<b>3 LIMITATIONS</b>	<b>4</b>
<b>4 RESULTS</b>	<b>4</b>
<b>5 CONCLUSION</b>	<b>7</b>

---

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 titles in the series. If you have questions or suggestions on any of these documents, please send them to [octo@icann.org](mailto:octo@icann.org).

---

# Executive Summary

ICANN's New Generic Top-Level Domains (gTLD) Program: 2026 Round will use a method to evaluate the likelihood of name collision risk for new gTLDs being added to the global DNS. To develop this method, OCTO analyzed the query traffic captured by the root servers. The study described in this paper is meant to help understand the similarities and differences in ranking of TLDs observed at the various root servers.

The primary purpose of our research is to determine whether the rankings seen by a single root server can be used to accurately represent all sets of root servers, or whether more root servers need to be included to get a meaningful sample. The result of this research shows that a single root server indeed could be representative.

## 1 Introduction

When estimating domain popularity, simple measurements such as the volume of queries do not necessarily reflect the popularity of a domain. The number of distinct networks that send DNS requests, rather than just the total number of requests, better reflects the overall popularity of the domain's services.

[DNS Magnitude](#), used in this study, is a popularity ranking metric for domain names first introduced by the [Austrian ccTLD registry](#). It uses statistically rigorous mathematics to combine query volume with the number of networks the queries come from into a metric that is more reflective of popularity than volume alone is.

A large fraction of the queries in this study are for domains under existing TLDs. Another large fraction is for undelegated TLDs. While some of those queries are for names under TLDs that may be added to the global DNS in the future, others are for domains that will never appear in the global DNS root zone.

## 2 Methodology

DNS Magnitude ranking has been calculated for traffic observed at L-root – the root server operated by ICANN – daily since 2021. The [current rankings](#) are available in a table that can be sorted and selected by different criteria. Only TLDs that conform to the hostname rules in [RFC 1035](#) are included: they must start with a letter, end with a letter or digit, and have only letters, digits, and hyphens as interior characters. In addition, the length of the label must be at least two characters to be included in the table. The table includes both labels that are delegated in the root of the global DNS and those that are not.

Traffic is observed over a single 24-hour period starting at midnight UTC. All queried names are aggregated by their TLD and source addresses are aggregated by /24 (for IPv4) and /48 (for IPv6) networks. The magnitude value is the natural logarithm of the unique sources observed for a top-level domain, divided by the natural logarithm of all unique sources, multiplied by 10 to obtain a value between 1 and 10. Higher values indicate queries originating from a larger diversity of networks. The rankings are shown only for the top 2,000 highest-ranked TLDs.

---

A root server operator (RSO) operates a root server identifier (RSI). RSIs are typically referred to by the letter that is associated with their full host name in the root zone of the global DNS. For example, the RSI managed by ICANN is typically called “L-root”, the RSI managed by Netnod is typically called “I-root”, and so on. More detail on RSOs and their respective letters can be found at [root-servers.org](https://root-servers.org).

To compare DNS magnitude ranking for data from L-root to data from other RSIs, OCTO used traffic captures to calculate DNS Magnitude ranking for those other RSIs. The "Day in the Life" (DITL) traffic capture project is a large-scale DNS-traffic data-collection event coordinated by [DNS-OARC](https://dns-oarc.org). Various DNS operators (particularly root server operators) capture DNS traffic over the same fixed period (generally 1-2 days) so that researchers can examine real-world DNS behavior. DNS-OARC's DITL traffic for 10 April, 2024, was used to calculate DNS Magnitude scores.

To understand equivalence between RSIs, we compute rank correlations between the RSIs using [Spearman's  \$\rho\$](https://en.wikipedia.org/wiki/Spearman%27s_rank_correlation_coefficient)  (Spearman's rank correlation coefficient) over the top 2,000 TLD lists for each RSI. Spearman's  $\rho$  measures how similarly any two lists are ordered. Instead of comparing actual values, it simply compares rank positions. If two rankings put the same TLDs near the top and bottom, the  $\rho$  value is high (closer to 1); if they are opposite, it is low (closer to -1). It is a straightforward calculation, does not assume any special data properties, and directly shows how closely two rankings agree. We calculate this for each combination of two RSIs and provide the results in a heatmap that shows the strength of correlation between two RSIs.

Another measure of equivalence is the [Jaccard index](https://en.wikipedia.org/wiki/Jaccard_index). Instead of looking at the ranking, it measures the extent to which two sets are similar by comparing what they have in common to what they have in total. If both sets of 2,000 TLDs are identical, the value is 1. If say, 1,500 TLDs are observed in both sets, the value is 0.75. A value of 0 means they share nothing. It is useful for comparing lists when you want to measure how common TLDs are, independent of where they're ranked. We calculated this metric for both the top 500 and top 2,000 TLDs.

## 3 Limitations

The DITL data used in this measurement covers a single UTC day. Equivalence across RSIs on that day may not generalize across time or deployment changes.

During assessment of available traffic data, there were issues with the data available for some of the RSIs. For E-root, the DITL capture covers only a single anycast instance out of the 328 in the RSI, skewing the results. For F-root, the dataset includes traffic from only 83 instances of 355 total in the RSI. For I-root, source IP addresses are anonymized, and this anonymization interferes with IP address aggregation and skews the results. Because of these inconsistencies, the DITL data from these three RSIs (E-root, F-root, and I-root) has been removed from the results below.

## 4 Results

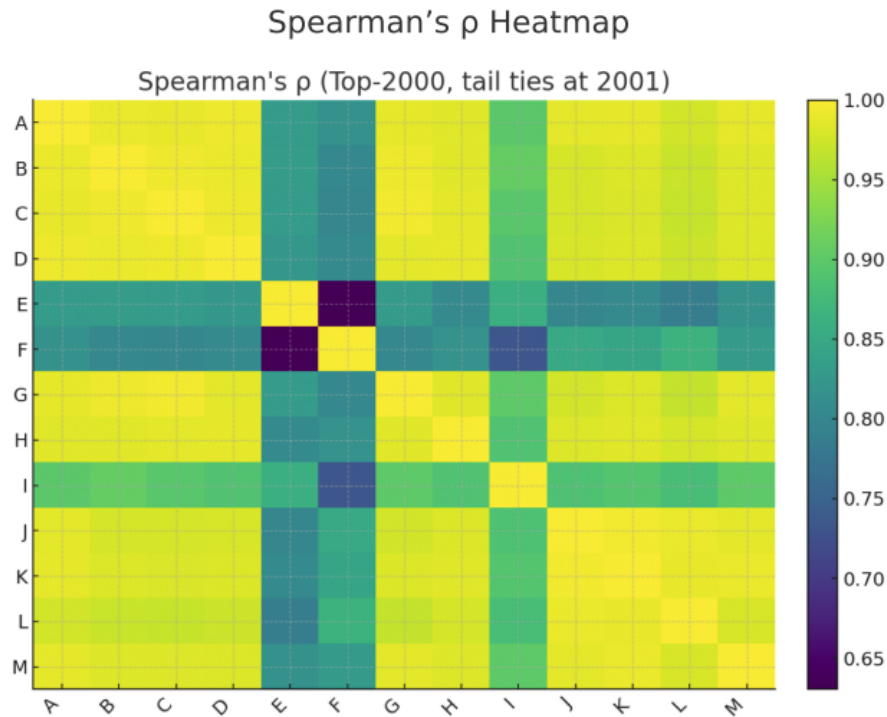
Spearman's  $\rho$  shows a strong correlation between the DNS Magnitude ranking of the top 2,000 TLDs between RSIs. The lowest observed value is 0.968 between G-root and L-root, meaning the two RSI's rank orders are almost identical across the full top 2,000. While differences exist,

they're minuscule and mostly in the fine ordering rather than wholesale rank shifts. The following matrix shows the values for each RSI pair.

	A	B	C	D	G	H	J	K	L	M
A	1.000	0.991	0.990	0.991	0.988	0.986	0.986	0.987	0.976	0.988
B	0.991	1.000	0.993	0.990	992	0.985	0.978	0.981	0.970	0.984
C	0.990	0.993	1.000	0.992	0.995	0.987	0.977	0.981	0.969	0.983
D	0.991	0.990	0.992	1.000	0.987	0.988	0.980	0.982	0.972	0.982
G	0.988	0.992	0.995	0.987	1.000	0.984	0.976	0.982	0.968	0.987
H	0.986	0.985	0.987	0.988	0.984	1.000	0.983	0.985	0.978	0.983
J	0.986	0.978	0.977	0.980	0.976	0.983	1.000	0.994	0.991	0.987
K	0.987	0.981	0.981	0.982	0.982	0.985	0.994	1.000	0.989	0.991
L	0.976	0.970	0.969	0.972	0.968	0.978	0.991	0.989	1.000	0.978
M	0.988	0.984	0.983	0.982	0.987	0.983	0.987	0.991	0.978	1.000

Table 1. Correlation between ranking of RSIs.

This table can be shown as a heatmap as well, where the value is expressed in a range between dark green and yellow. The following heatmap shows that values for the RSIs not included in Table 1 are outliers.



Graph 1. Heatmap showing correlation between RSIs.

To assess the commonality of TLDs in the lists, we calculated the overlap for the first 500 TLDs and the first 2,000 TLDs using the Jaccard overlap coefficient. By focusing on the first 500 of the lists, we emphasize how common the top 500 TLDs are, which implies how consistent the highest-ranked TLDs are across different RSIs.

	A	B	C	D	G	H	J	K	L	M
A	1.000	0.988	0.984	0.988	0.980	0.931	0.984	0.976	0.961	0.980
B	0.988	1.000	0.980	0.980	0.980	0.934	0.972	0.972	0.953	0.976
C	0.984	0.980	1.000	0.980	0.992	0.934	0.984	0.980	0.969	0.976
D	0.988	0.980	0.980	1.000	0.984	0.927	0.988	0.980	0.965	0.980
G	0.980	0.980	0.992	0.984	1.000	0.931	0.976	0.980	0.961	0.976
H	0.931	0.934	0.934	0.927	0.931	1.000	0.927	0.927	0.912	0.927
J	0.984	0.972	0.984	0.988	0.976	0.927	1.000	0.984	0.976	0.976
K	0.976	0.972	0.980	0.980	0.980	0.927	0.984	1.000	0.972	0.980
L	0.961	0.953	0.969	0.965	0.961	0.912	0.976	0.972	1.000	0.953
M	0.980	0.976	0.976	0.980	0.976	0.927	0.976	0.980	0.953	1.000

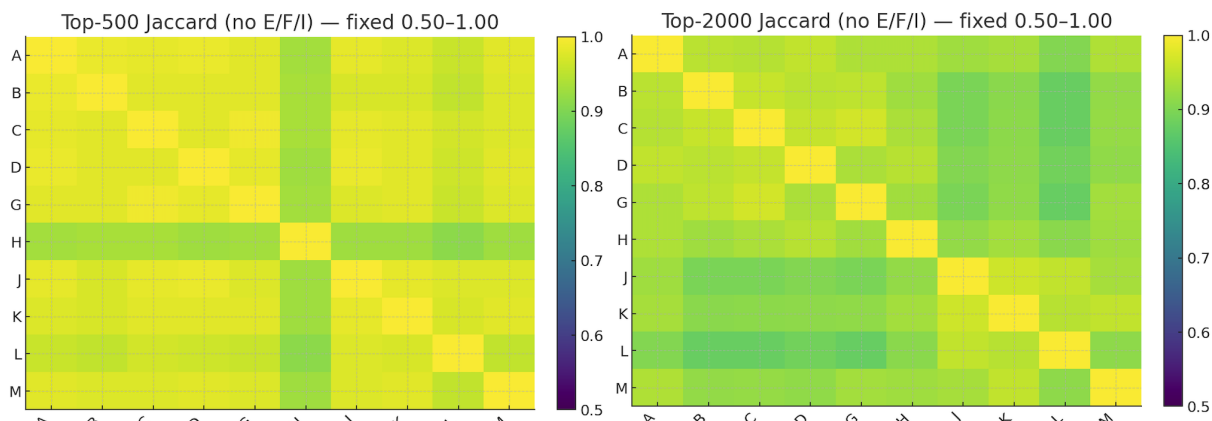
Table 2. Top-500 overlap between RSIs.

The lowest value in Table 2 is 0.912 between H-root and L-root. An overlap of 0.912 is still high. It means the two sets being compared share about 91.2% of their TLDs. In the context of top 500s for the various RSIs, this indicates a significant overlap: the vast majority of TLDs in the top of one list also appear in the top of the other. In other words, the top of the list of ranking is highly consistent across two RSIs.

	A	B	C	D	G	H	J	K	L	M
A	1.000	0.948	0.944	0.954	0.939	0.938	0.926	0.933	0.904	0.940
B	0.948	1.000	0.958	0.947	0.952	0.928	0.896	0.908	0.878	0.917
C	0.944	0.958	1.000	0.956	0.968	0.937	0.895	0.912	0.879	0.918
D	0.954	0.947	0.956	1.000	0.936	0.946	0.904	0.914	0.889	0.916
G	0.939	0.952	0.968	0.936	1.000	0.930	0.896	0.915	0.875	0.931
H	0.938	0.928	0.937	0.946	0.930	1.000	0.918	0.931	0.908	0.927
J	0.926	0.896	0.895	0.904	0.896	0.918	1.000	0.964	0.955	0.933
K	0.933	0.908	0.912	0.914	0.915	0.931	0.964	1.000	0.946	0.953
L	0.904	0.878	0.879	0.889	0.875	0.908	0.955	0.946	1.000	0.914
M	0.940	0.917	0.918	0.916	0.931	0.927	0.933	0.953	0.914	1.000

Table 3. Top 2,000 overlap between RSIs.

The lowest value in Table 3 is 0.875 between G-root and L-root. An overlap of 0.875 is high. It means the two sets being compared share about 87.5% of their TLDs. In the context of top 2,000 overlaps for the various RSIs, this also indicates a significant overlap: the vast majority of TLDs in the top of one list also appear in the top of the other. This ranking is highly consistent in both the top 500 and the top 2,000 TLDs.



Graph 2. Heatmap illustration of the two Jaccard matrices for the top 500 and top 2,000 TLDs. The range of colors is mapped on values between 0.5 and 1.0.

## 5 Conclusion

The results of this DNS Magnitude root server equivalence study demonstrate a remarkably high degree of consistency across the analyzed RSIs. Despite operational, geographical, and architectural diversity among the RSOs and their respective RSIs, the correlation of DNS Magnitude rankings is extremely strong, with Spearman’s  $\rho$  values consistently exceeding 0.97 for the top 2,000 TLDs. This indicates that the rank ordering of TLDs based on DNS Magnitude is nearly identical across independent RSIs.

Jaccard overlap coefficients for the top 500 and top 2,000 TLDs confirm this. Even the lowest observed Jaccard values (0.912 and 0.875, respectively) reflect a high level of overlap, meaning that over 87% of the most popular TLDs appear consistently across RSIs. Such strong agreement implies that DNS Magnitude measurements are highly consistent, independent of which RSI’s traffic is analyzed.

Outliers such as E-root, F-root, and I-root can be explained by dataset limitations, such as restricted capture scope, incomplete anycast coverage, and anonymization effects, rather than intrinsic inconsistencies in DNS query traffic. When these outliers are excluded, the equivalence across the remaining RSIs is statistically near complete.

This study concludes that the DNS Magnitude ranking derived from a single, well-represented RSI provides a representative reflection of global DNS query diversity at all RSIs. For most research and operational purposes, analyzing one RSI is sufficient to approximate the aggregate view of all, enabling efficient and scalable measurement without sacrificing accuracy.