

# Differences in Monitoring the DNS Root Over IPv4 and IPv6

Tarang Saluja<sup>1</sup>, John Heidemann<sup>2</sup>, Yuri Pradkin<sup>2</sup>

<sup>1</sup>Swarthmore College; Swarthmore, PA

<sup>2</sup>University of Southern California Information Sciences Institute; Marina Del Rey, CA

**Abstract**—The Domain Name System (DNS) is an essential service for the Internet which maps host names to IP addresses. The DNS Root Server System operates the top of this namespace. RIPE Atlas observes DNS from more than 11k vantage points (VPs) around the world, reporting the reliability of the DNS Root Server System in DNSmon. DNSmon shows that loss rates for queries to the DNS Root are nearly 10% for IPv6, much higher than the approximately 2% loss seen for IPv4. Although IPv6 is “new,” as an operational protocol available to a third of Internet users, it ought to be just as reliable as IPv4. We examine this difference at a finer granularity by investigating loss at individual VPs. We confirm that specific VPs are the source of this difference and identify two root causes: VP *islands* with routing problems at the edge which leave them unable to access IPv6 outside their LAN, and VP *peninsulas* which indicate routing problems in the core of the network. These problems account for most of the loss and nearly all of the difference between IPv4 and IPv6 query loss rates. Islands account for most of the loss (half of IPv4 failures and 5/6ths of IPv6 failures), and we suggest these measurement devices should be filtered out to get a more accurate picture of loss rates. Peninsulas account for the main differences between root identifiers, suggesting routing disagreements root operators need to address. We believe that filtering out both of these known problems provides a better measure of underlying network anomalies and loss and will result in more actionable alerts.

## I. INTRODUCTION

Automatic monitoring is essential to ensuring that production systems operate 24/7. A number of commercial and non-commercial monitoring systems exist today and provide minute-by-minute updates about many public and private network services.

The Domain Name System (DNS) is a distributed directory for the Internet [9]. DNS maps hostnames to IP addresses so that users can use names like [example.com](http://example.com) while packets are sent to 192.0.2.1 (it provides other information as well). DNS names are hierarchical, and above the top-level domains ([.com](http://.com), [.edu](http://.edu), [.jp](http://.jp), etc.) is the DNS root. The DNS root is provided by the Root Server System [15], a group of 12 organizations that operate 13 services. Because DNS is part of every web request and e-mail message, its operation is critical.

RIPE Atlas is a measurement system with more than 11k vantage points (VPs, called RIPE Atlas Probes) spread around 171 countries globally [13], [14]. Operated as a service of RIPE, it allows third-party measurements to be taken at no cost. RIPE’s DNSmon [12] uses Atlas to monitor the DNS Root Server System, and the DNSmon dashboard provides both a near-real time and historic view of Root DNS reliability.

DNSmon’s dashboard suggests that IPv6 access to the DNS Root is consistently worse than IPv4. Access to the same




Figure 1. DNSmon on 2022-07-23 with the color scale set from 3% (green) to 8% (red).

13 services often has 3–6% failure rates over IPv6, while IPv4 usually sees no failures. For example, Figure 1 shows DNSmon for all of 2022-07-23. IPv4 is every first line, and each unit of time is green, indicating 0–3% loss during that unit of time. IPv6 is every second line, and each unit of time is orange or red, indicating 3–8% loss during that unit of time. RIPE examined their data and found that some Atlas VPs were misconfigured—although they removed these VPs from reporting, the problem returned.

The goal of our work is to evaluate RIPE Atlas data over time to understand the causes for IPv4 and IPv6 differences. Is the discrepancy between IPv4 and IPv6 simply a measurement error, or are there underlying differences between the protocols which indicate that IPv6 is less reliable? How often do these measurement errors or underlying differences occur?

Our work makes three contributions. First, we classify Atlas VP connectivity by using the idea of *islands* and *peninsulas* [3]. This helps us differentiate the root-causes of reachability problems local to the VP from those in the core of the Internet (§II). This work builds on prior evaluation of failed VPs by RIPE (see §V), but we identify both islands and peninsulas rather than just islands. Second, we show that

the brunt of the difference between IPv4 and IPv6 is due to problems with the measurement system (misconfigured VPs), not in the protocols or services. In particular, most of the difference is due to VPs that are IPv6 islands—VPs that are misconfigured and cut off from the IPv6 network ([§III-D](#)). Removing these misconfigured VPs greatly reduces differences in reliability ([Figure 10](#)). Third, after accounting for configured VPs, we show that remaining differences between root DNS services can be identified as peninsulas ([§III-E](#)) and explained as routing disagreements and other partial connectivity issues ([Figure 11](#)). Finally, we show that these results are hold over many days, although the particular VPs with problems change slowly ([§IV](#)). Together, our observations provide a diagnosis of the difference between IPv4 and IPv6 reliability and provide solutions for how IPv6 routing can be improved and brought nearly on-par with IPv4 reliability.

We hope that RIPE Atlas operators can use our tools to automatically tune DNSmon and provide a more accurate reflection of general IPv6 reliability. We also hope that our work can motivate the networking community to identify and improve IPv6 routing to provide connectivity as seamless as IPv4.

Our analysis is based on public data from RIPE Atlas. We will make our analysis code and results available publicly on our website [[16](#)].

## II. APPROACH

### A. Problem Statement

By examining RIPE Atlas observations, we identify when observers see problems and suggest potential root-causes for these problems. Using information about these root causes, we aim to provide measurements which best reflect what is experienced by the typical user of DNS. Furthermore, we hope to give the operators of RIPE Atlas and the DNS Root Server System information which can be used to distinguish between measurement problems and actual routing problems.

The target of our measurements is the *DNS Root Server System* [[15](#)], composed of 13 independent systems identified by *identifiers* A through M. All identifiers commit to serve the same data, the DNS root at the top of the domain name systems' hierarchy (above [.com](#)). Each identifier provides service on a unique IPv4 and IPv6 address, and for capacity and reliability, each address uses IP anycast [[10](#)], [[1](#)] and is provisioned by computers in multiple locations. [Figure 2](#) shows the 1563 locations of the DNS root server system on a map, as of 2022-07-18.

We use data recorded by the RIPE Atlas measurement system. Atlas has more than 11k vantage points (VPs, called RIPE Atlas Probes) spread around 171 countries globally [[13](#)], [[14](#)]. [Figure 3](#) shows the VPs that were part of RIPE as of 2022-07-13.

RIPE Atlas takes measurements of the DNS Root Server System frequently and reports that data in DNSmon [[12](#)]. By default, DNSmon uses data from about 100 RIPE Atlas *anchors*. Anchors are VPs which are well connected (often in data centers). For our analysis, we also examine data from all RIPE Atlas VPs, using all anchors and Atlas probes that report data from measurements.




Figure 2. The DNS Root Server System on 2022-07-18. Each dot is an “instance” offering service, with instances in the same location offset in circles (when run by different operators) or lines (when run by the same operator).



Figure 3. RIPE Atlas Vantage Points on 2022-08-02.

Our primary dataset is 24 hours of data covering all of 2022-07-23 (UTC). We selected this day at random when we began our work. On this day, we see data from 10,082 VPs, with 10,082 and 5,173 claiming IPv4 and IPv6 connectivity respectively. In [§III-C](#), we will show that some VPs claim connectivity but apparently cannot reach outside their local-area network.

We confirm that our results on this day are typical by comparing them to data taken over the 7 subsequent days, from 2022-07-20 to 2022-07-26 (inclusive), as discussed in [§IV](#).

All RIPE Atlas data is publicly available [[11](#)], as are daily reports with our analysis [[17](#)]. We will provide our analysis scripts [[16](#)], allowing anyone to reproduce our results.

Although our specific quantitative results concern the Root Server System as measured from RIPE Atlas, we expect these results will apply to similar large services using IP anycast, and could use other measurement systems. Such services include Content Delivery Networks and DNS services by Akamai, Amazon, Cloudflare, Google, and Microsoft. Measurement systems similar to RIPE Atlas include CAIDA's Archipelago [[5](#)] and several commercial systems.

### B. Data Collection

Each RIPE Atlas VP queries each of the 13 Root identifiers over both IPv4 and IPv6 with several different types of queries: DNS SOA requests, ICMP echo requests, and traceroutes. VPs query at different rates based on user request and VP type. About 869 VPs are designated as anchors, which are well-connected servers that are often located at data centers. These VP anchors typically send queries to each identifier every 5 minutes. The majority of VPs are small, embedded devices,

root identifier	anchors		all VPs	
	IPv4	IPv6	IPv4	IPv6
A	1423335	1423343	10009	11009
B	10100811	8912024	10010	11010
C	1423315	1615204	10011	11011
D	1423383	1423391	10012	11012
E	1423399	5118710	10013	11013
F	1423323	1423339	10004	11004
G	1423375	6918264	10014	11014
H	3082611	3082616	10015	11015
I	1423363	1423355	10005	11005
J	1423367	1423347	10016	11016
K	1423387	1423331	10001	11001
L	1423371	3636450	10008	11008
M	1423327	1423379	10006	11006
period	300 s	300 s	240 s	240 s
Active VPs	25	52	10,082	5,173

Table I

RIPE ATLAS MEASUREMENT IDs FOR SOA MEASUREMENTS FROM ANCHORS AND ALL VPs BY PROTOCOL (IPv4 AND IPv6), AND HOW OFTEN EACH VP QUERIES AND ALSO NUMBER OF ACTIVE VPs ON 2022-07-23.

often deployed in homes. These also probe each root, but typically every 4 minutes. Other query types (like traceroute) are done less frequently or on demand. In all cases, VPs can be tasked by researchers to make additional queries.

Since we compare IPv4 and IPv6 to examine their reliability, we look at data from all VPs and for both IPv4 and IPv6.

**Table I** lists the specific RIPE Atlas measurement IDs we draw upon. All RIPE data is available for public download with these IDs [14], so others can reproduce our results.

We reproduced the DNSmon results using SOA queries from anchors. We find that SOA records from all VPs provide a similar result. As there are more observers, using all VPs provides better precision. We also confirm that measurements using ICMP Echo Request show similar reliability.

Our analysis in this paper uses SOA queries from all VPs. For all of our analysis we download a 24 hour set of RIPE Atlas data in JSON format using their APIs.

### C. Detecting Problems: Islands and Peninsulas

Our goal is to distinguish measurement problems from routing problems by examining Atlas VPs. We now describe how to detect the root-causes we look for in VP data.

**Candidate Root-Causes:** We look for two problems: *edge routing problems*, where a VP thinks it has IPv6 support but cannot actually reach any of the IPv6 network, and *core routing problems*, where a VP has IPv6 support and can reach some of the IPv6 Internet but not all of it.

Both of these problems are harmful to the users, who are unable to rely on their IPv6 connections, but they have different solutions. Edge routing problems must be solved at the edge, by the VP operator, while core routing problems stem from routing interactions between ISPs. These core routing problems should be brought to the attention of network operators for consideration in their choices of routing and peering.

**Islands and Peninsulas Detect These Problems:** To understand these problems and how to detect them, we use recent work examining partial Internet connectivity [3]. That work

defines islands as VPs which cannot reach any of the Internet, and peninsulas as VPs which can reach part of the Internet but not all of it. We use these concepts in our algorithms to detect islands and peninsulas in the RIPE data. Islands suggest edge routing problems, and peninsulas suggest core routing problems.

For each VP, we consider 24 hours of queries to each root identifier. Typically each (VP, identifier)-measurement is performed every five minutes, yielding 288 observations over the day (VP failures or reboots can reduce this count). As we report in §III-A, we find that for the majority of VPs (about 90%), either all queries succeed, or all queries fail. In the few remaining cases, either a few queries intermittently fail or block(s) of queries fail.

To simplify our analysis, we classify (VP, identifier) combinations into “always fail” or “did not always fail” over the 24h period. We use always-fails to identify root causes. Root-cause identification is used to detect some VPs as faulty due to measurement errors. Eliminating measurement errors causes remaining problems to stand out, making it easier for operators to distinguish between problems that need attention and problems in the measurement system.

*Islands* are VPs that cannot reach anywhere in IPv6 over the 24h. That is, an island is a VP where (VP, \*) is “always fail” for all identifiers.

*Peninsulas* are VPs that can reach some identifiers, but never reach other identifiers over the day. That is, a peninsula is a VP where there exists at least one identifier  $\alpha$  for which (VP,  $\alpha$ ) is “always fail” and at least one identifier  $\beta$  where (VP,  $\beta$ ) is “does not always fail”.

**Quantifying Reliability:** We characterize typical failure rates for all queries (§III-B) in IPv4 and IPv6. We then use these definitions to look at the failure rate for subsets of VPs after we remove islands and then both islands and peninsulas respectively. These values allow us to compare what DNSmon reports for the two network protocols IPv4 and IPv6. However, unlike DNSmon, we can do these comparisons for all VPs rather than a small selection of anchors.

If we can correct different types of routing failures, we can then study what the internet would be like without these failures. In §III-D, we consider failure rates after discarding islands, since islands indicate measurement errors due to misconfigured VPs that have edge routing problems. Then in §III-E, we consider failure rates after discarding both islands and peninsulas to understand what the underlying IPv4 and IPv6 reliability is like.

We suggest that our reports about islands and peninsulas can guide operators. Islands require the attention of RIPE Atlas operators (or operators of specific Atlas VPs) as islands are VPs which are misconfigured and should be corrected or removed from service. Peninsulas require the attention of ISPs and Root operators, since peninsulas are indicative of routing problems in the network core. Some peninsulas may involve multiple parties, making ruling out islands and confirming peninsula stability important.

## III. EVALUATION OF ONE DAY




Figure 4. A cumulative distribution of query loss per VP for IPv6 B root SOA queries on 2022-07-23.

#### A. How Often Do Queries Fail?

We first consider how often queries from VPs fail. Our analysis considers only VPs where all queries fail (§II-C). If many VPs show partial failures (only fail some of the time), this choice would lead us to ignore an important part of the data.

We classify VPs by what fraction of their queries fail. **Figure 4** shows a cumulative distribution function of the fraction query failures for all VPs. The data for the cumulative distribution function is SOA queries over IPv6 from about 5k VPs to B-Root (RIPE Measurement ID 8912024). We identify 4 groups, from left to right: *always succeed* (green, with 0% loss), *rarely fail* (light yellow, with some loss, but less than 5%), *often fail* (orange, with more than 5% but less than complete loss), and *always fail* (red, 100% loss).

By far the largest category is always succeed—around 80% of VPs never see failures, confirming that networks generally work without loss. Second largest is always-fail, accounting for about 10% of IPv6 VPs. We use these two categories in our analysis. We focus on identifying always-fail VPs for two reasons. First, they represent a serious and persistent problem, since users at the same location as the VP will never see service. Second, they indicate *actionable* problems. In particular, since they persist, there is time for an operator to investigate the VP, confirm its persistent failure, and then take some action to correct the problem.

The 10% of remaining probes are split between rare and frequent loss. We do not consider these VPs in our analysis because they are less important and more difficult to resolve quickly. Rare loss is consistent with packet loss due to occasional congestion that is expected in a best-effort network [6], and occasional loss from transient routing changes [18].

VPs which often fail are a more serious problem. Manual examination shows that VPs which often fail have persistent outages that last minutes or hours. Although medium-term failures are difficult for users and worthy of attention, we




Figure 5. Stacked bars of the four VP categories by protocol (IPv4, left, blue; IPv6, right, red), for each identifier. Dataset: 24 h on 2022-07-23.

ignore these in our current analysis on the assumption that this state is transient and the VP will join the set of VPs with either occasional failure or always failure soon. If the VP had a series of failures that were resolved, it does not need operator attention and we should ignore it. If the frequent failure was the start of a long-term routing problem, the VP will show up as always-fail on the next day and get our attention then.

If our work is successful in helping resolve persistent problems, then reexamining frequent or occasional failures will become more important. It is possible that frequent or occasional failures are not transient, but rather indicate a pattern of failure or a more fundamental routing problem.

**Figure 5** shows the classifications for each identifier in both IPv4 and IPv6. We see that the fraction of always-fails is similar across all identifiers for both IPv4 and IPv6, although IPv6 consistently has more always-fails than IPv4, as we will explore in §III-B.

#### B. Does IPv6 See More Loss Than IPv4?

We next compare the overall failure rate of protocols to answer “is IPv6 worse than IPv4”?

We evaluate the fraction of lost queries for all queries from all VPs (about 10k reporting IPv4 and 5k IPv6, **Table I**) to each root identifier, reporting the mean fraction of loss for each hour over 24 hours in **Figure 6**. Each pair of bars compare IPv4 (left, blue) and IPv6 (right, red) loss rates for each of the 13 root identifiers. **As reported by RIPE Atlas, IPv6 loss is consistently much higher than IPv4: IPv6 is around 10% while IPv4 is around 2%.**

This result confirms what DNSmon reports in **Figure 1**: IPv6 is “more red”. However, DNSmon is based on data from 50–100 RIPE anchors, while **Figure 6** uses all available VPs, about 100× more.

Small black error bands on each bar show  $\pm 2$  standard deviations around each mean, providing roughly a 95% confidence interval. These very small error bands show that hourly loss rates are very consistent over the day. We confirm this




Figure 6. Query loss fractions for each root identifier for IPv4 (left) and IPv6 (right) for all available Atlas VPs. Dataset: 24 h on 2022-07-23.




Figure 7. Fraction of VPs by potential problem. Dataset: 24 h on 2022-07-23.

result in [Figure 10](#), where the solid lines compare hourly loss rates for just B-Root on the same day.

### C. Where are These Problems?

As measured, IPv6 shows much more loss than IPv4. This difference is puzzling. Even though IPv6 is a much newer protocol, it has been in wide operation for more than a decade. It should be just as reliable as IPv4.

To understand this difference, we next identify potential root causes by classifying VPs into islands (cannot reach any identifier) and peninsulas (can reach some identifiers but not others). [Figure 7](#) shows the fraction of VPs in each class. In the middle we also include protocol-exclusive islands; these islands contain VPs for one protocol only.

**Islands:** We first consider islands: about 8% of VPs are IPv6 islands (400 of 5173) while only about 2% of VPs are IPv4 islands (205 of 10,082). As such, the proportion of VPs which have local IPv6 misconfigurations is four times more than the proportion of VPs which have local IPv4 misconfigurations. VPs with IPv6 misconfigurations are likely on an IPv6 LAN which cannot route IPv6 to the rest of the world.

To understand if VPs have completely misconfigured networking or if the misconfiguration is protocol specific, we




Figure 8. Fraction of all VPs that are peninsulas for each root identifier. Dataset: 24 h on 2022-07-23.

found that there were 32 VPs which were islands in both IPv4 and IPv6. The middle columns (“exclusive islands”) compare islands after these VPs are discarded. Given that it causes such a small difference, we treat those 32 VPs like all others.

**Peninsulas:** Finally, we show the fraction of VPs in peninsulas with the right bars of [Figure 7](#). We see very slightly more peninsulas than islands for IPv4 and very slightly less for IPv6, with 8% of VPs seeing IPv6 peninsulas (396 of 5173) and 2% IPv4 (239 of 10,083). Again, IPv6 has about four times more peninsula problems than IPv4. Since a VP observing peninsulas cannot also be an island, these problems are in addition to VPs that are islands. Furthermore, while islands affect the fraction of query success for all roots equally, peninsulas affect the query success of some roots and not others.

Since peninsulas indicate the inability to reach some identifiers, we can compare which identifiers have the most problems with connectivity. [Figure 8](#) shows what fraction of VPs cannot reach a given root identifier but can reach at least one other root identifier. Note that some VPs cannot reach multiple identifiers, but each VP that is a peninsula must be able to reach some identifier. Hence, the same VP can be represented in multiple bars of [Figure 8](#), and the bars are, by definition, smaller than the total number of peninsulas in [Figure 7](#).

We see some variation in which identifiers are least reachable, with most seeing peninsula reachability problems for less than 1% of VPs. However, nearly 5% of VPs are peninsulas which cannot reach C-Root. That is by far the largest proportion of VPs which cannot reach a identifier due to being part of a peninsula. We believe this trend is because of a known routing dispute between Cogent (operator of C-Root) and Hurricane Electric [\[8\]](#). Since both provide IPv6 service to many users, a lack of IPv6 routing between them shows up as peninsulas in RIPE Atlas. Resolution of this problem depends on reaching business agreements to exchange traffic directly or through a third party. In addition, several other identifiers show many more IPv6 peninsulas than IPv4.




Figure 9. Comparing query loss fractions for each root identifier for IPv4 (left) and IPv6 (right) as measured from all VPs except those identified as islands. Dataset: 24 h on 2022-07-23.

#### D. Fixing Observation Problems by Removing Islands

In our discussion of root-causes in §II-C, we suggested that a common cause of an island is a misconfiguration at a VP which allows it to think it is on the Internet but leaves it unable to route globally. Such VPs are not suitable for a measurement system—they should be excluded from consideration. Our expectation is that removing misconfigured observers will make it easier for us to detect and resolve other Internet problems.

To produce a corrected measurement system, we detect islands and remove island VPs as described in §II-C. Figure 9 shows the revised evaluation of query reliability (compare it to Figure 6 with all VPs). **If we do not consider islands, we measure that IPv4 fails 1/2 as often and IPv6 fails 1/6th as often—most RIPE Atlas failures are problems local to the VP, not at the DNS root server**

Figure 9 shows IPv4 queries fail with fractions of approximately 0.005 to 0.01, and IPv6 from approximately 0.01 to 0.06. We also see that IPv6 is much closer to IPv4 than before: the IPv6 failure rate is about double IPv4, as opposed to four times IPv4.

Results are again stable over each hour of our 24 h measurement period. The dashed lines in Figure 10 show that failure fractions at B-Root are stable each hour over the day for VPs omitting islands, and it confirms that the failure fractions are much lower than with islands and IPv6 is double IPv4.

However, some identifiers still show lower IPv6 connectivity than IPv4. The IPv4/v6 difference is largest for C-root. Removing islands amplifies the relative differences due to routing problems in the network core, as described previously in §III-C.

#### E. Observing Core Routing Problems and Peninsulas

We next consider routing problems in the Internet core. We detect such problems as peninsulas. That is when a VP is partially functioning because it can see some root identifiers, but it has persistent inability to connect with others. Peninsulas represent meaningful problems that will be visible to end-users




Figure 10. IPv4 and IPv6 query failure rates for each hour to B-Root. Solid lines show data for all VPs, Dashed lines omit VPs which are islands. Dataset: 24 h on 2022-07-23.




Figure 11. IPv4 and IPv6 query failure rates for each hour to B-Root. Dashed lines omit VPs which are islands. Dotted lines omit VPs which are islands or peninsulas. Dataset: 24 h on 2022-07-23.

as broken websites. Such problems cannot be easily addressed by users.

Our ability to identify peninsulas means we can carry out a thought-experiment: what would the Internet be like if these detectable routing problems were resolved? Would that address the remaining differences between IPv4 and IPv6? We can study this question by excluding VPs that are peninsulas from observation.

Figure 12 shows query failure fractions after VPs that are either islands or peninsulas have been removed (compare it to Figure 6 and Figure 9).

**If we remove peninsulas and islands, all root identifiers have similar, quite small loss fractions. After addressing VP problems (islands), persistent routing failures (peninsulas) affecting some root identifiers account for most of differences in loss rates.** In Figure 12, loss is between 0.004 to 0.009 for IPv4 and 0.008 to 0.016 for IPv6. This analysis




Figure 12. Comparing query loss fractions for each root identifier for IPv4 (left) and IPv6 (right) as measured from all VPs except those identified as islands or peninsulas. Dataset: 24 h on 2022-07-23.

also confirms that the large differences between IPv4 and IPv6 are due to peninsulas.

We see similar results when we look specifically at B-Root over all 24h in [Figure 11](#), a version of [Figure 10](#) with a rescaled  $y$  axis. For individual hourly observations, without islands and peninsulas, we now begin to see more variation across the day. This variation is likely because smaller effects are now visible, such as random loss due to network congestion.

Although observations with neither islands nor peninsulas show lower loss rates, IPv6 still sees about twice the loss rate of IPv4. This difference suggests that IPv4 does currently have a lower end-to-end packet loss rate, at least for DNS queries.

#### IV. EVALUATION OF MULTIPLE DAYS

Our evaluation ([§III](#)) is based on data from one day, 2022-07-23 (a Saturday). We next show data for that full week, from 2022-07-20 (three days before) to 2022-07-26 (three days after). **This evaluation shows that the most VPs that see islands and peninsulas experience them for many days, but about one-twentieth change each day.** The stable core suggests that operators of RIPE Atlas and the DNS Root identifiers have time to diagnose and address problems.

##### A. Is the Fraction of Query Failures Stable Over Time?

We examine IPv4 and IPv6 loss rates in [§III-B](#) and show that IPv6 loss is much higher than IPv4 loss. While loss is stable for a given identifier over 24 h, there is some variation across the identifiers.

To confirm that our previous results are true in general, we reexamine this data by adding days around our target day (2022-07-23). We processed 3 days before and 3 days after, the week from 2022-07-20 to 2022-07-26. [Figure 13](#) shows the same data as [Figure 6](#), but with each bar replaced by seven bars for each individual day of the week.

Comparing each group of seven bars, we see that query loss rates are fairly stable over the course of a week, confirming that our results in [Figure 6](#) are representative.

While there are some instances where the error bands are larger than expected (for example, both protocols for B-Root on 2022-07-22, and both protocols for K-Root on 2022-07-21), this occurs rarely and without a discernible pattern. The day after a day of high variability often shows a lower-than-typical reliability, suggesting variability is due to transient routing issues that become a persistent problem, as discussed in [§III-A](#)

##### B. Are the Fractions Of Islands and Peninsulas Stable?

We identified islands and peninsulas as problems VPs encounter in [§III-C](#) and showed that more VPs see such problems in IPv6 than IPv4 based on one day of data. We next reexamine that data over a week to confirm that day was typical.

[Figure 14](#) shows what fraction of VPs are islands (to all identifiers) and peninsulas (to at least one identifier) for each day of the week from 2022-07-20 to 2022-07-26. As with query rates, we see the fraction of VPs reporting each problem is fairly consistent across all seven days. (As a large measurement system, RIPE VPs fail and recover, so the exact number of VPs reporting on each day varies by a small amount: 3–36 of about 10k for IPv4 and 3–13 of about 5100 VPs for IPv6.)

##### C. Are Islands VPs Stable or Transient?

While we have shown in ([§IV-B](#)) that the fraction of islands is stable over a week, we do not yet know if *specific* VPs which are islands *remain* islands. In particular, it is possible that the set of VPs which are islands changes substantially even though the fraction of islands is stable. It is important to understand whether or not *specific* VPs are stable because debugging routing problems requires operator-to-operator communication (opening trouble tickets and exchanging email). If the *specific* VPs which are islands keeps changing, then mitigation will be difficult because operators will be less clear about which VPs need to be fixed.

In [Figure 15](#), we show how often the same islands appear in a stacked bar chart, and then changes on adjacent days for IPv6. The middle bar shows the fraction of islands on our primary day, 2022-07-23. The bars to the left and right then compare one day earlier or later with that day, showing island VPs in common with the 23rd in blue and new island VPs (those that were not islands on the 23rd) in orange (on top). The two bars on the ends do the same thing for two days before and after the 23rd, showing overlap with the 23rd in blue and with the new VPs on the 22nd or 24th in orange, and new island VPs on this day in green.

Given the large blue section on each day, we see that islands VPs are quite stable. There is very little churn – only about 1 in 20 islands VPs are new on each day, meaning that approximately 19 in 20 island VPs stay that way from one the day to the next. In fact, if we look at all seven days, we can see that approximately 85–90% islands on any given day are also islands for the other seven days of the week for both IPv4 and IPv6.




Figure 13. Query loss fractions for each day over one week, for each root identifier, for IPv4 (left, blue bars in each group) and IPv6 (right, red bars). Dataset: 7 days, from 2022-07-20 to 2022-07-26.




Figure 14. Fraction of VPs by potential problem, for each day over one week, by protocol (IPv4 is light blue, left group, IPv6 in pink on the right). Dataset: 7 days, from 2022-07-20 to 2022-07-26.




Figure 15. Overlap of island-VPs for two days before and after 2022-07-23, for IPv6.




Figure 16. Overlap of peninsula VPs for two days before and after 2022-07-23, for IPv6.

#### D. Are Peninsula VPs Stable or Transient?

We just showed that island VPs are stable over a week (§IV-C). We next repeat this analysis for peninsula VPs.

In Figure 16, we see how often the same peninsulas appear in a stacked bar chart. This can be interpreted like Figure 15, as described near the end of §IV-C. Again, we see that that peninsula VPs are quite stable, with minimal churn. Although there is lower incidence of orange and green in this stacked bar chart, that is because it appears that 7/23 has more peninsulas than neighboring days. If 7/24 was our day of reference, then we would see plenty of orange on 7/23.

Given the large blue section on each day, we see that peninsula VPs are quite stable. There is very little churn – only about 1 in 40 peninsula VPs are new on each day, meaning that approximately 39 in 40 peninsula VPs stay that way from one the day to the next. In fact, if we look at all seven days, we can see that approximately 90% peninsulas on any given day are also peninsulas for the other seven days of the week for both IPv6. Similarly, for IPv4, 75% of peninsulas on any given day are also peninsulas for the other seven days of the week. In general, peninsulas have similar performance to islands, with only IPv4 peninsulas having a bit less stability than other categories.

#### E. Summarizing Stability and Long-term Operation

Overall, these results show that while there is some churn (VPs that change status on different days), the majority of VPs that are islands or peninsulas persist for at least a week. We believe that islands reflect VP network misconfigurations and peninsulas show problems with routing in the core network. Both phenomena are properties of network configuration (at home or by ISP routing policies), and so this stability is consistent with this interpretation.

This stability suggests that operators have time to diagnose and correct routing problems. The high level of consistency over multiple days confirms the persistence and relevance of these problems.

However, there is much churn and that highlights the importance of periodic monitoring of the measurements. New islands and peninsulas often occur, so vigilance and maintenance is required to resolve routing problems. The resolution of these routing problems is necessary to make DNSmon more accurate. Increased sensitivity will allow it to detect more subtle, transient problems, like congestion. The small, but non-trivial, variation justifies the need to continuously monitor for changes in peninsula VPs.

**Operation:** To support diagnosis of these problems, we have automated operation of these algorithms and provide a daily report of islands and peninsulas for RIPE Atlas evaluation of the Root DNS [17].

## V. RELATED WORK

Our work builds on several areas of prior work.

Our work directly builds upon reports by RIPE Atlas operators that there are VPs which claim to support IPv6 but cannot. While they periodically examine the set of VPs used in DNSmon for islands, our work shows that regressions to

islands seem fairly common among VPs. In addition, while RIPE Atlas operators focused only on islands, we also examine the effect of peninsulas.

We also build on an analysis of partial Internet connectivity and the terms islands and peninsulas [3]. That work also examined islands and peninsulas in RIPE Atlas. However, while it focused only on IPv4, we also examine IPv6.

Bush et al. explored partial connectivity in the control-plane (routing with BGP), comparing it to data-plane reachability [4]. They show control-plane reachability can overestimate. Accordingly, we focus only on data-plane reachability.

Several systems have recognized partial connectivity (peninsulas) and proposed to address it in an overlay network [2] or by route manipulation [7]. We instead suggest that islands are errors that should be addressed in the measurement system by ignoring such VPs, and we do not propose specific solutions to peninsulas.

## VI. FUTURE WORK

In this paper, we addressed root causes associated with VPs which have total failure with respect to a root server (no successful queries for 24 hours). However, we did not look at VPs which have more ambiguous performance. These VPs are responsible for the remaining loss and hence the remaining discrepancy.

To diagnose the root causes for the VPs which have some failures, we must better understand whether these VPs are experiencing failure *randomly* or *sequentially*. While *random* failures are interspersed throughout the observation interval with no apparent patterns, *sequential* failures are consecutive, occurring one after another at certain times during the same observation window. *Sequential failure* is more likely due to temporary islands and peninsulas formed by short-term outages, re-routing events, and network congestion. On the other hand, *random failure* is more likely associated with random packet drops, not necessarily caused by congestion.

We can differentiate between *sequential failure* and *random failure* by using the CUSUM statistical test to check for randomness. This method can help develop a heuristic to identify more *temporary* islands and peninsulas. Such future work can help locate root causes for the discrepancies which still exists after removing islands and peninsulas which exist at the granularity of at least a day.

## VII. CONCLUSION

In this paper, we identify two root causes (islands and peninsulas) for the discrepancy in query failure fraction between IPv4 and IPv6, and we further analyze their impact on the discrepancy. Our analysis reveals that while much of the discrepancy is caused by faulty measurement from islands, partial routing failure of peninsulas also contributes to the discrepancy. Furthermore, we find that even if we ignore both islands and peninsulas, there is still a discrepancy which should be investigated further. We hope that our code and analysis can be used by root name server and RIPE Atlas operators to filter misleading data appropriately, fix faulty measurements from islands, and find or shed light on

solutions for peninsulas. Finally, we identify future work, on VPs which fail some of the time, which can help elucidate the discrepancy that continues to exist after the removal of islands and peninsulas.

#### ACKNOWLEDGMENTS

Tarang Saluja began this work during the summer of 2022, under the USC/ISI Research Experience for Undergraduates program (NSF grant 2051101, PI: Jelena Mirkovich). The work of Yuri Pradkin and John Heidemann is supported in part by the project “CNS Core: Small: Event Identification and Evaluation of Internet Outages (EIEIO)” (CNS-2007106). We thank Guillermo Baltra for his work on island and peninsula detection [3] and sharing his results.

We very much thank RIPE, and the many hosts of RIPE VPs, for providing Atlas as a open platform to the community, We also thank several root operators for their feedback on this work.

#### REFERENCES

- [1] J. Abley and K. Lindqvist. Operation of anycast services. RFC 4786, Internet Request For Comments, December 2006. (also Internet BCP-126).
- [2] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. Resilient overlay networks. In *Proceedings of the Symposium on Operating Systems Principles*, pages 131–145, Chateau Lake Louise, Alberta, Canada, October 2001. ACM.
- [3] Guillermo Baltra and John Heidemann. What is the internet? (considering partial connectivity). Technical Report arXiv:2107.11439v2, USC/Information Sciences Institute, May 2021.
- [4] Randy Bush, Olaf Maennel, Matthew Roughan, and Steve Uhlig. Internet optometry: assessing the broken glasses in Internet reachability. In *Proceedings of the ACM Internet Measurement Conference*, pages 242–253. ACM, November 2009.
- [5] CAIDA. Archipelago (Ark) measurement infrastructure. website <https://www.caida.org/projects/ark/>, 2007.
- [6] David D. Clark. The design philosophy of the DARPA Internet protocols. In *Proceedings of the 1988 Symposium on Communications Architectures and Protocols*, pages 106–114. ACM, August 1988.
- [7] Ethan Katz-Bassett, Colin Scott, David R. Choffnes, Ítalo Cunha, Vytautas Valancius, Nick Feamster, Harsha V. Madhyastha, Tom Anderson, and Arvind Krishnamurthy. LIFEGUARD: Practical repair of persistent route failures. In *Proceedings of the ACM SIGCOMM Conference*, pages 395–406, Helsinki, Finland, August 2012. ACM.
- [8] Rich Miller. Peering disputes migrate to IPv6. website <https://www.datacenterknowledge.com/archives/2009/10/22/peering-disputes-migrate-to-ipv6>, Oct. 22 2009.
- [9] P. Mockapetris. Domain names—concepts and facilities. RFC 882, Internet Request For Comments, November 1983.
- [10] C. Partridge, T. Mendez, and W. Milliken. Host anycasting service. RFC 1546, Internet Request For Comments, November 1993.
- [11] RIPE NCC. RIPE Atlas. web site <https://atlas.ripe.net/>, 2010.
- [12] RIPE NCC. DNSMON. web site <https://atlas.ripe.net/dnsmon/>, 2015.
- [13] RIPE NCC Staff. RIPE Atlas: A global Internet measurement network. *The Internet Protocol Journal*, 18(3):2–26, September 2015.
- [14] RIPE NCC Staff. Welcome to RIPE Atlas! website <https://atlas.ripe.net/>, July 2022.
- [15] Root Operators. <http://www.root-servers.org>, April 2016.
- [16] Tarang Saluja. Finding and fixing persistent problems in IPv6 monitoring of DNS. <https://ant.isi.edu/reu/2022/>, August 2022.
- [17] Tarang Saluja and Yuri Pradkin. RIPE Atlas islands and peninsulas. [https://ant.isi.edu/ripe\\_atlas\\_islands/](https://ant.isi.edu/ripe_atlas_islands/), September 2022.
- [18] Feng Wang, Zhuoqing Morley Mao, Jia Wang, Lixin Gao, and Randy Bush. A measurement study on the impact of routing events on end-to-end Internet path performance. In *Proceedings of the ACM SIGCOMM Conference*, pages 375–386, Pisa, Italy, August 2006. ACM.