A Primer in Registration Data Access Protocol (RDAP) Performance

ICANN Office of the Chief Technology Officer

Carlos Gañán
OCTO-024
17 May 2021

ICANN

TABLE OF CONTENTS

EXECUTIVE SUMMARY 4

1 INTRODUCTION 4

2 METHODOLOGY 6

2.1 Creating a Sample of Queries 6
2.1.1 RDAP Queries for TLD Registries 6
2.1.2 RDAP Queries for RIRs 7
2.1.3 RDAP Queries for ICANN-accredited Registrars 8
2.2 Running the Queries 8

3 RDAP SERVICES 9

3.1 RIRs' RDAP Services 10
3.2 TLD Registries' RDAP Services 10
3.3 ICANN-Accredited Registrars' RDAP Services 10

4 THE PERFORMANCE METRIC 10

5 AGGREGATED RESULTS 11

5.1 Response Time 11
5.2 Breakdown of HTTPS Transaction Timings 12
5.3 Response Time Per Vantage Point 12
5.4 RDAP Response Size 13
5.5 Response Time Versus the Source IP Address Type 14

6 RIRS' RDAP RESPONSE TIME 15

6.1 Response Time Versus the Vantage Point 16
6.2 Response Time Versus the Source IP Address Type 16
6.3 Response Time Versus the Response Size 17

7 TLD REGISTRIES' RDAP RESPONSE TIME 18

7.1 Response Time Versus the Vantage Point 19
7.2 Response Time Versus the Source IP Address Type 19
7.3 Response Time Versus the Response Size 20

8 ICANN-ACCREDITED REGISTRARS' RDAP RESPONSE TIME 21

8.1 Response Time Versus the Vantage Point 22
8.2 Response Time Versus the Source IP Address Type 22
8.3 Response Time Versus the Response Size 23

9 LIMITATIONS **24**

10 CONCLUSIONS **24**

A REGISTRARS' RDAP ERRORS **25**

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 improve the shared responsibility for upholding the security and stability of the DNS by strengthening DNS coordination in partnership with relevant stakeholders. 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 Domain Name System (DNS) is essential to the overall functioning of the Internet. With millions of new domain names registered every year, it is important to have a mechanism that enables timely access to accurate details describing who has registered each domain. WHOIS has been the main communication protocol to query registration data for more than thirty-five years, however it presents several limitations such as lack of standardized query and response formats and inability to authenticate users.

To overcome these limitations, in 2015, the Internet Engineering Task Force (IETF) proposed a new protocol standardizing registration data access while supporting internationalized domain names (IDNs) and registration data, in addition to client authorization, and other features. Known as the Registration Data Access Protocol (RDAP), this new protocol enables access to registration data for current domain names, IP addresses, and Autonomous System Numbers (ASNs). This study investigates the performance of the RDAP services that have been deployed as it is an ICANN requirement that generic top-level domain (gTLD) registries and ICANN-accredited registrars must implement an RDAP service by the end of August 2019.

As part of this study, during December 10, 2020 to January 10, 2021, over seven million RDAP domain queries were executed from ten different vantage points against RDAP services operated by four of the five regional Internet registries (RIRs), all gTLD registries, the ccTLD registries that have deployed RDAP, and all ICANN-accredited registrars. The measurements show that the average query response time was 1.02 seconds; however there are significant differences depending on the RDAP operator. It was observed that, on average, the RIRs provide the fastest RDAP service followed by gTLD registries, followed by ccTLDs, followed by ICANN-accredited registrars.

After performing an exploratory statistical analysis of the RDAP query response time, several factors that impact the response time were identified:

- *Query origin*: The source location of the query had a significant impact on the query response time. Queries originating from Europe and North America received, on average, faster responses than those from Asia or Africa.
- *Source IP Address Type:* Queries executed over IPv6 had, on average, lower response times than those over IPv4.
- *RDAP response size:* despite RDAP responses varying from a few hundred to hundreds of thousands of bytes, response size did not have a statistically significant impact on the response time.
- *Secure channel*: as RDAP leverages a secure communication protocol that provides encryption, it introduces additional overhead compared to WHOIS. It was observed that the time required to create a secure channel with the RDAP server accounted, on average, for forty percent of the total response time.

1 Introduction

Internet identifiers — namely Internet Protocol (IP) addresses and domain names — are essential for the functioning of nearly all Internet protocols, e.g., routing protocols such as the Border Gateway Protocol (BGP) or application layer protocols such as the Hypertext Transfer Protocol (HTTP). These identifiers are managed by a group of interdependent organizations that assign and allocate identifiers, from the five Regional Internet Registries (RIRs) that manage Internet number identifiers to the thousands of domain name registries and registrars that manage the registration of Internet domain names.

To be able to identify who is responsible for a given identifier at any point in time, in the early 1980s, the WHOIS protocol was specified to provide access to registration information. Even though the original design of WHOIS presented several limitations such as the lack of a standardized response format and an inability to authenticate users, it became the main mechanism for external parties to obtain registration data about Internet identifiers. This registration information is useful not only for network administrators to fix system problems, but also to combat threats such as spam or fraud. Thus, timely and accurate registration data is important and provides useful information for mitigating potential threats.

In 2015, the Internet Engineering Task Force (IETF) proposed a new protocol to standardize registration data access while supporting internationalized domain names (IDNs) and client authorization. It was envisioned that it would replace the WHOIS protocol. Known as the Registration Data Access Protocol (RDAP), this new protocol enables access to registration data for current domain names, IP addresses, and Autonomous System Numbers (ASNs). By the end of August 2019, generic top-level domain (gTLD) registries and ICANN-accredited registrars were required by ICANN to implement RDAP.

This paper investigates the performance of the RDAP service that has been deployed in response to the ICANN requirement that gTLD registries and registrars must implement an RDAP service by 26 August 2019. Previous research by Interisle and by Viagenie examined whether RDAP service operators comply with ICANN's policies and deployment issues.

The purpose of this study is to assess the performance of deployed RDAP services. It uses remote active measurements to estimate the response time for accessing registration data of currently deployed RDAP services and identifies potential bottlenecks. Response time is a critical metric as it determines the amount of time that any client has to wait to obtain the response back. While the response time of RDAP is expected to be higher than the WHOIS protocol due primarily to the overhead introduced by a TLS handshake, the authors are unaware of any prior study that examines how much time it takes to get a registration record via RDAP.

For this study, the authors focused on three different types of operators involved in the deployment of RDAP services: top-level domain (TLD) registries, ICANN-accredited domain name registrars, and the RIRs for comparison purposes only. This study measures the response times and validity of the responses for the RDAP services provided by each of these types of operators.

The rest of this report is structured as follows. Section 2 presents the methodology to measure the efficiency of RDAP. Section 3 describes the current deployment status of RDAP for each provider. Section 4 defines the metric used to measure the efficiency of each RDAP service. Section 5 presents the aggregated results followed by Sections 6, 7, and 8, which present the results for RIRs, TLD registries and ICANN-accredited registrars, respectively. Section 9 states the limitations of this study, and Section 10 gives conclusions from the study. Appendix A lists specific registrars' RDAP services that did not reply to queries.

2 Methodology

A measurement infrastructure was deployed to estimate different metrics related to the performance of RDAP services. All measurements were conducted remotely by an RDAP client simply executing test RDAP queries against the RDAP servers. Because the probes were external to the RDAP server, metrics related to the internal usage of resources such as CPU or memory usage were not quantified. The method is quite simple and consists of 2 steps: (i) creating a random sample of domain queries, and (ii) executing the queries while recording the response and response time.

2.1 Creating a Sample of Queries

In order to test the response time of each RDAP service, it is necessary to create a sample list of queries. While some RDAP services allow queries for different types of resources (mainly IP address, autonomous system numbers, or domains), this study focused on domain queries as this is the object type that is common to all RDAP services including those provided by RIRs, registries, and registrars. To capture potential differences among the responses of the same RDAP service, 100 different queries were created for each service. This threshold was chosen heuristically to make the measurement process viable in the client machines with limited resources.

Three different processes are followed in order to create RDAP queries depending on the type of RDAP operator, i.e., RIRs, TLD registries, and ICANN-accredited registrars.

2.1.1 RDAP Queries for TLD Registries

Acquiring the samples can be visualized as follows:

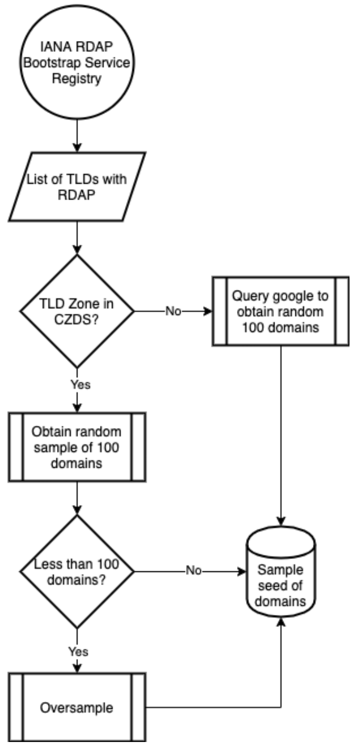

To create this sample list of domains, the authors performed the following steps:

1. Obtained a list of TLDs for which Internet Assigned Numbers Authority (IANA) has an RDAP base URL listed in the bootstrap service.
2. For each TLD:
    a. If the zone file was available, e.g., via the Centralized Zone Data System (CZDS), the authors obtained a random set of 100 domains. If the zones did not contain 100 domains, then the authors oversampled the domains;
    b. If the zone file was not available, the authors obtained a random set of 100 unique domains by extracting second-level domains (SLDs) from the results of a web search (e.g., Google query: “ * site: ‘TLD’ “). If the results of the search did not contain 100 different domains, then the authors oversampled the domains.
3. Created domain queries by appending the sample set of domains to each RDAP URL as stated in RFC 7482.

2.1.2 RDAP Queries for RIRs

1. Obtained the list of regional RIR RDAP servers from IANA’s bootstrap list.
2. For each RIR RDAP server, the authors created a set of 100 domains using the in-addr.arpa public zone files as the RIR RDAP servers only provide registration data for the in-addr.arpa domains they are delegated.
3. Created domain queries by appending the sample set of domains to each RDAP URL from step (1) following RFC 7482.

2.1.3 RDAP Queries for ICANN-accredited Registrars

1. Obtained the list of base URLs for the RDAP servers operated by registrars from IANA's bootstrap list.
2. For each domain in the .com zone file, obtained the corresponding registrar via the WHOIS protocol.
3. Obtained a sample of 100 domains for each registrar that had an RDAP server mapped in step (1) by using the registrar information of step (2)
    a. If there were not 100 .com domains, then the authors oversampled the domains.
    b. If there were no .com domains for a particular registrar, then the authors used the Spamhaus pDNS API to retrieve domains for that registrar by appending its name to the URL "https://api.spamhaus.net/dbl/v1.2/registrar/".
4. Created domain queries by appending the sample set of domains to each RDAP URL from step (1) following RFC 7482.

2.2 Running the Queries

To measure the response time of the different RDAP services, ten different virtual machines (VM) were provisioned to conduct the measurements. Each virtual machine was located in a different autonomous system spread across six different continents to account for potential measurement biases due to routing. Figure 2.1 shows the location of these VMs. Table 2.1 specifies the network prefixes of these VMs along with the city and country where they were located.

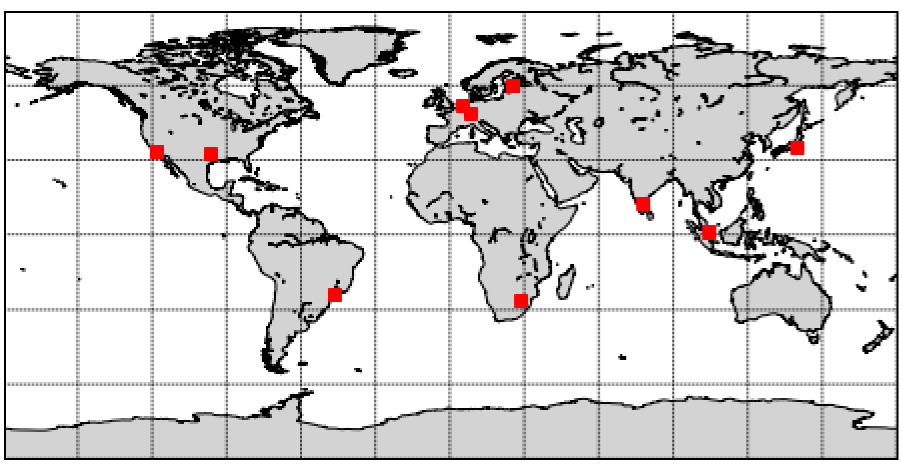

Figure 2.1. Geolocation of the measurement VMs (red dots)

Prefix	ISO 3166 Country Alpha2-code	City
104.156.236.0/23	US	Texas
145.220.0.0/16	NL	Utrecht

208.77.188.0/22	US	California
35.247.192.0/18	BR	São Paulo
172.104.160.0/19	SG	Singapore
13.244.0.0/15	ZA	Cape Town
218.227.0.0/16	JP	Tokyo
51.210.0.0/16	FR	Paris
135.181.0.0/16	FI	Helsinki
206.189.128.0/20	IN	Bengaluru

Table 2.1. Prefix and location of the measurement VMs

The queries ran periodically, e.g., every 5 minutes, and were executed at the same time from the different nodes. However, each RDAP service was only queried once every 5 minutes by each node. This is necessary to distinguish the RDAP services that share the same infrastructure. The RDAP services were grouped based on the IP address and the origin AS of their domains.

Next, the RDAP queries were executed as follows:

1. Every 5 mins, one random query for each one of the RDAP services was selected.
2. The queries were executed and the response times were recorded.

The measurements are aimed at quantifying a response time. The response time is one of the most important metrics to track the performance of a REST API, because it captures the quality-of-service as experienced by the users. For the purpose of this study, response time is defined as the time elapsed since an RDAP query is executed until the response is received minus the time to resolve the domain name.

3 RDAP Services

During this investigation, the RDAP services provided by RIRs, TLD registries, and ICANN-accredited registrars are examined. The first but essential step before measuring the performance of these services is to identify where they are running. IANA provides the lists of the base URLs of these services as part of the RDAP bootstrap service.

In registries that operate multiple TLDs, some of those TLDs share the same RDAP URL. Thus, the number of URLs below does equal the total number of TLDs in the DNS.

3.1 RIRs' RDAP Services

The base URLs of the RIRs' services are maintained by IANA and are published here. Each individual URL corresponds to one of four RIRs. Latin America and Caribbean Network Information Centre's (LACNIC) RDAP service did not allow querying for domains at the time this study was designed, and thus it was not measured during this study.

3.2 TLD Registries' RDAP Services

A total of 818 RDAP base URLs are specified by IANA in the TLD bootstrap service as of December 2020. However, these RDAP base URLs are not unique, and thus, it is reasonable to assume that it is the same server providing the RDAP service for multiple TLDs. RDAP URLs whose domain names resolve to IP addresses belonging to the same AS are grouped together to avoid being queried more than once from the same vantage point during a measuring period. In total, 39 different ASes were identified; 28 of these services provide registration data for more than one TLD.

3.3 ICANN-Accredited Registrars' RDAP Services

A similar process was followed to obtain a list of base URLs for the RDAP services provided by ICANN-accredited registrars. A total of 350 unique RDAP base URLs for registrars were specified by IANA's bootstrap list as of December 2020. Nevertheless, not all these URLs were run by independent service providers. Hence, the RDAP services are aggregated by AS as in the case of gTLDs. In total, 168 different ASes were identified.

4 The Performance Metric

Using the timing information provided by the curl tool, the following metric is defined:

- **Response time**: the time from the start of the request until the last byte is received minus the time to resolve the RDAP service domain name.

This time is computed by adding up the following times:

- **Connect time:** the time it took from the start of the connection until the connection to the remote host (or proxy) was completed.
- **Appconnect time:** the time it took from the end of the connection time until the TLS handshake with the remote host was completed.
- **Pretransfer time:** the time it took from the end of the TLS handshake until the registration data transfer is just about to begin.
- **Start-transfer time:** the time it took from the end of the pretransfer until the first byte is received.
- **Transfer time:** the time it took from the first byte until the last byte is received.

Note, that some services redirect the queries to other URLs. In this case, the "response time" metric will also account for these redirections by calculating the total time from query to response including any redirection.

5 Aggregated Results

After measuring the RDAP response time over four weeks from December 10, 2020 until January 10, 2021, a total of 7,353,811 valid measurements were collected. Around 0.1% of the measurements were discarded due to errors in the response.

5.1 Response Time

It was observed that, on average, RIRs provide the fastest RDAP service followed by TLD registries, followed by ICANN-accredited registrars. Of all queries, 95% were answered in less than 4 seconds. Table 5.1 evidences the presence of outliers in the measurements. In the case of the RDAP response times provided by registrars, the standard deviation is 150% higher than the mean value. Looking at the maximum values, it is also worth noting that some queries took more than 15 minutes to get a response. These were extreme outliers as shown by the 99th percentile.

Response time (sec)

	mean	std	min	50%	95%	99%	max
RIR	0.88	1.03	0.02	0.66	1.94	3.15	132.20
TLD	1.26	1.38	0.04	1.06	2.82	14.44	259.37
Registrar	1.46	2.47	0.03	1.01	3.76	11.93	940.41

Table 5.1 Summary statistics of the RDAP response time per operator type

Figure 5.1 shows the distribution of the response time per RDAP operator type. The significant differences can be seen when observing the size of the boxes. In this figure, outliers have been removed from the boxplots; outliers here are defined as data points beyond 1.5 times the interquartile range from the box quartiles.

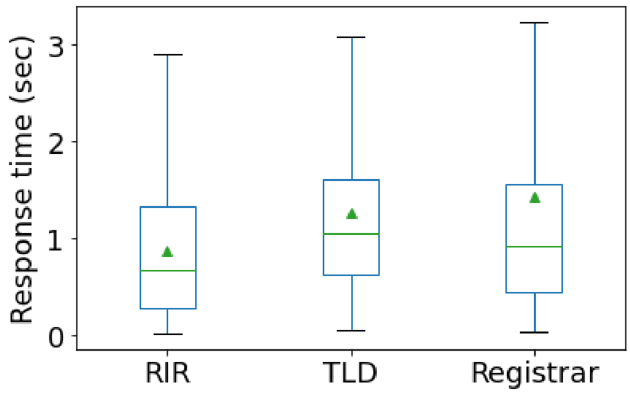

Figure 5.1. Box plots of the response time per RDAP operator type

5.2 Breakdown of HTTPS Transaction Timings

All the queries in this study were performed using HTTPS; thus the total response time is also affected by the TLS handshake. To understand the extent to which the HTTPS protocol affects RDAP performance, the total response time is broken down into the five metrics, which are explained in Section 4.

Table 5.2 shows the mean value of the different timings. Between 320 to 430 milliseconds were spent on average per RDAP query on the TLS handshake, and between 226 to 500 milliseconds were spent on transferring the response data.

	Time (sec)				
	Connect time	**Appconnect time**	**Pretransfer time**	**Starttransfer time**	**Transfer time**
RIR	0.13	0.32	0.00	0.20	0.06
TLD	0.19	0.43	0.00	0.40	0.04
Registrar	0.19	0.42	0.00	0.49	0.01

Table 5.2 Breakdown of the different HTTP transaction timings

As can be seen in Figure 5.2, on average, the TLS handshake takes around 40% of the total response time. Most of the time (~60%) is spent in starting the connection, and transferring the RDAP response only accounts for less than 40% of the total response time. (The pretransfer times are 0 and thus not visible in Figure 5.2.)

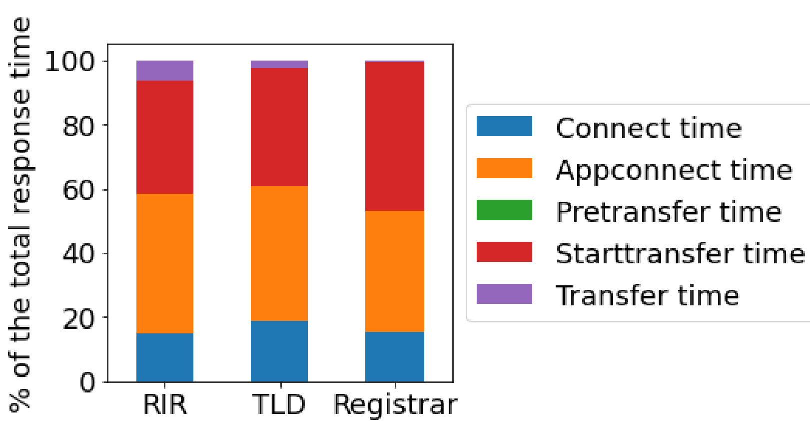

Figure 5.2 Breakdown of RDAP response time

5.3 Response Time Per Vantage Point

The response time per query was also measured for each of the ten vantage points deployed across the globe. Significant differences were observed depending on the location from which the queries were run. Table 5.3 shows the main descriptive statistics for the response time per vantage point. The response time measurements were significantly higher from the vantage points located in India, Japan, Singapore, and Brazil.

Response time (sec)

location	mean	std	min	50%	95%	99%	max
BR	1.69	1.71	0.07	1.34	3.40	5.34	144.15
FI	1.10	1.57	0.03	0.85	2.81	5.90	156.92
FR	1.11	1.80	0.02	0.82	2.66	6.49	153.32
IN	1.46	1.23	0.12	1.20	3.08	5.28	99.00
JP	1.54	1.43	0.07	1.28	3.17	5.86	94.45
NL	1.25	3.51	0.03	0.66	3.81	10.89	940.41
ZA	1.56	1.23	0.04	1.38	3.16	5.36	148.42
SG	1.59	2.53	0.05	1.26	3.27	9.46	199.57
US1	0.83	0.60	0.04	0.70	1.89	2.79	21.06
US2	1.13	1.76	0.03	0.77	2.50	5.81	165.67

Table 5.3. Response time depending on the location of the vantage point

The response time measurements suffer from a significant presence of outliers. This is represented by high standard deviation and maximum values that in some cases were 20 times larger than the 99th percentile.

Figure 5.3 shows the boxplots of response time distribution for each vantage point depending on the city where it is located. Again, the node in Brazil was the one that registered not only the highest average latency, but also the largest latency variance.

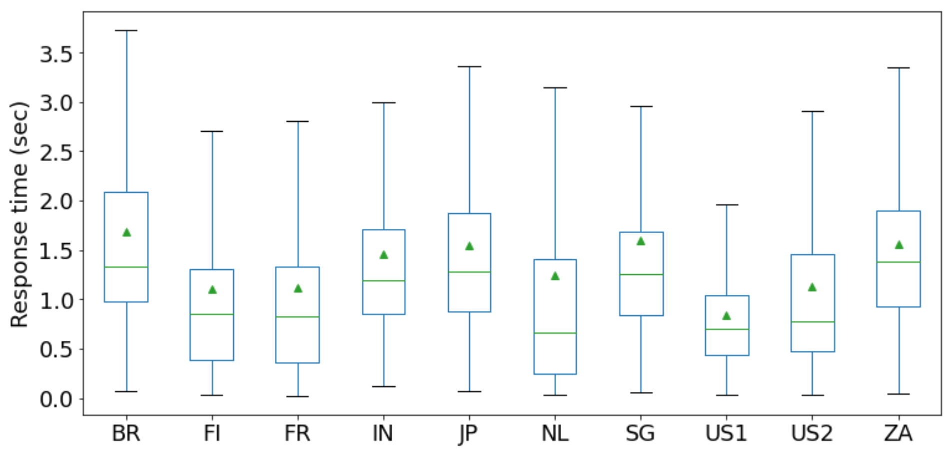

Figure 5.3. Response time versus the vantage point (no outliers; triangles showing average response time)

5.4 RDAP Response Size

Next, the response size in bytes for each RDAP query was calculated. The average response size of all the performed queries was around 6 kbytes with a standard deviation of 4.4 kbytes. This evidences a great diversity of response sizes across the different RDAP services that were queried. Table 5.4 breaks down the response size statistics per RIR, TLD registries, and ICANN-accredited registrars. As can be seen, registrars' RDAP services provided the shortest responses on average while some registrars provided a response of more than 122 kbytes.

Note, that the extremely large response size was due to a repeated field in the RDAP response vCard.

Response size (bytes)

	mean	std	min	50%	95%	max
RIR	6,157	5,189	1,059	3,307	16,557	28,546
TLD	6,931	4,730	664	6,257	11,280	61,241
Registrar	4,345	2,527	285	4,111	8,899	122,499

Table 5.4.Summary statistics for the RDAP response size

Finally, a potential relationship between the response time and size was investigated. Figure 5.4 plots these variables against each one of the performed queries. A Pearson product-moment correlation coefficient was computed to assess the relationship between the response size and response time. No strong correlation was found (Pearson r≈0, p-value≈0).

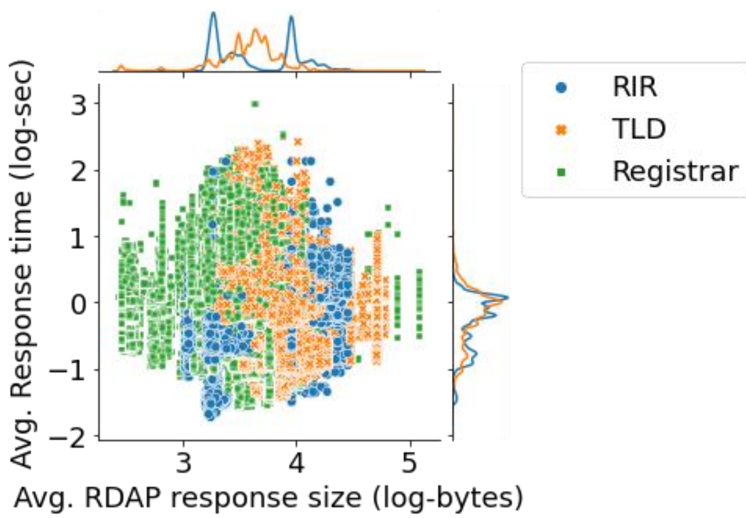

Figure 5.4.Response time versus the RDAP response size per operator type

5.5 Response Time Versus the Source IP Address Type

A subset of the vantage points nodes allowed the authors to perform RDAP queries using IPv6 source addresses. From these, the same IP queries were performed using IPv4 and IPv6 source addresses. As can be seen in Figure 5.5, the average response time for queries over IPv4 was 1.07 seconds while the same queries executed over IPv6 were responded slightly faster with an average response time of 0.87 seconds. Moreover, IPv4 queries suffered from more variance (std: 2.20 sec) than IPv6 queries (std: 0.74 sec).

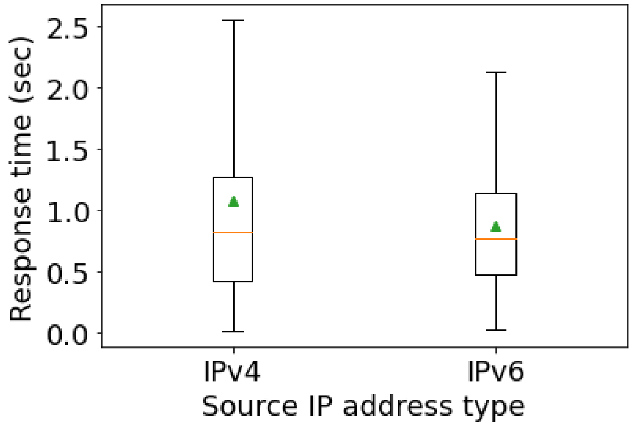

Figure 5.5 Response time per IP source address type (no outliers; triangles showing average response time)

6 RIRs' RDAP Response Time

Table 6.1 shows the main descriptive statistics of the response time for the RIRs' RDAP service. Note, that at the time data was collected, LACNIC did not provide domain information via RDAP. The African Network Information Centre's (AFRINIC) RDAP service has the slowest response time on average, and APNIC's RDAP service is the fastest. It is also worth noting that there were some outliers in the measurements. For instance, some queries took more than two minutes to return the corresponding RDAP response. Again, these were extreme outliers as 99% of all queries finished in less than four seconds.

	Response time (sec)						
RIR	**mean**	**std**	**min**	**50%**	**95%**	**99%**	**max**
AFRINIC	0.98	0.76	0.10	1.10	1.37	2.19	129.99
APNIC	0.33	0.89	0.07	0.25	0.72	1.13	92.65
ARIN	0.87	0.99	0.03	0.85	1.77	3.17	132.19
RIPE	0.42	0.41	0.02	0.50	0.69	0.87	65.08

Table 6.1 Summary statistics of the RDAP response time per RIR

Figure 6.1 shows the boxplot for these same measurements to better illustrate the differences in the response time distribution per RIR. As seen by the interquartiles, the RDAP service provided with the American Registry for Internet Numbers (ARIN) had the largest response time variance. APNIC research shows the smallest variance which can probably be attributed to their RDAP infrastructure being deployed in the cloud.

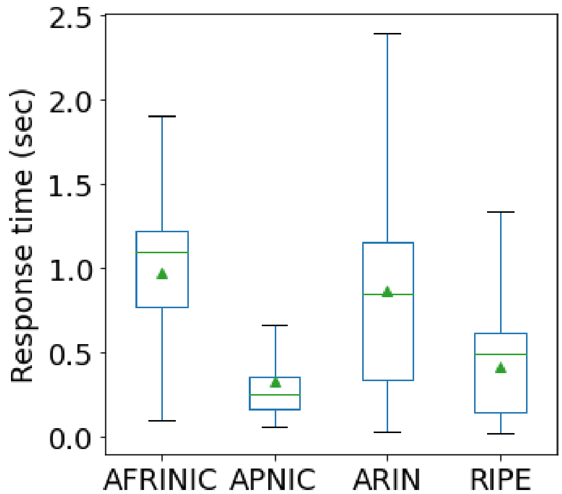

Figure 6.1. Box plots of the response time per RIR (outliers have been removed from the boxplots)

6.1 Response Time Versus the Vantage Point

The response time per query was also measured for each of the ten vantage points. Table 6.2 shows the summary statistics of these. Significant differences can be seen when comparing the response time for the different vantage points. For instance, the response time of the queries, which ran from the vantage points located in Brazil and India, were on average two times slower than the ones, which ran from the U.S. or the Netherlands.

	Response time (sec)						
location	**mean**	**std**	**min**	**50%**	**95%**	**99%**	**max**
BR	1.06	1.78	0.13	0.87	1.58	3.84	73.18
FI	0.59	0.62	0.03	0.60	1.15	1.78	92.19
FR	0.56	1.63	0.02	0.51	1.08	1.66	132.19
IN	1.04	0.62	0.13	1.17	1.81	2.94	26.06
JP	0.98	0.52	0.49	0.89	1.65	3.21	23.31
NL	0.53	0.50	0.03	0.60	0.95	1.63	23.14
ZA	0.85	0.97	0.10	0.63	1.78	2.46	92.47
SG	0.84	0.81	0.07	0.96	1.41	2.88	33.30
US1	0.53	0.45	0.10	0.34	1.26	1.31	21.05
US2	0.44	0.62	0.03	0.30	1.11	1.16	92.65

Table 6.2 RIRs' RDAP response time depending on the location of the vantage point

6.2 Response Time Versus the Source IP Address Type

A subset of the vantage point nodes was allowed to perform RDAP queries using IPv6 source addresses. From these, the same IP-queries were performed using IPv4 and IPv6 source addresses. As can be seen in Figure 6.3, no significant differences are found between the queries that used IPv4 compared to IPv6. The average response time for both was around 0.7

seconds. However, IPv4 queries suffered from more variance (std: 1.02 sec) than IPv6 queries (std: 0.68 sec).

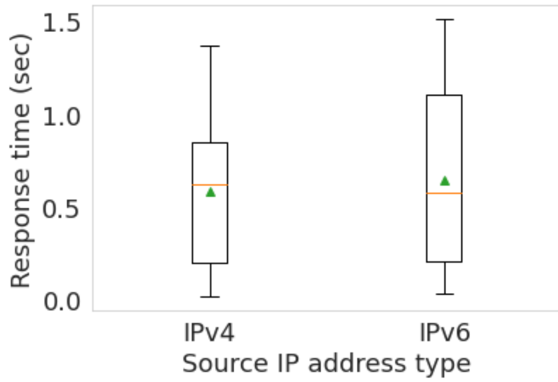

Figure 6.3 Response time per IP source address type (no outliers; triangles showing average response time)

6.3 Response Time Versus the Response Size

Finally, the impact of the RDAP response size on the response time is investigated. While the format of the RDAP response is standardized, the amount of information varies per RIR. Figure 6.4 shows the response size distribution per RIR. As expected, there are differences in terms of response size. ARIN's RDAP responses were 4 kbytes on average, while RIPE, AFRINIC and APNIC did not reach 3.5 kbytes on average.

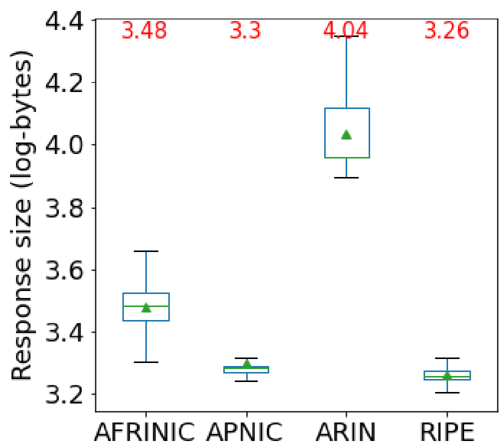

Figure 6.4 RDAP response size per RIR (no outliers; triangles showing average response time)

While the responses from ARIN's RDAP service were larger on average, the average response time was still faster than AFRINIC's RDAP service. Figure 6.3 shows the response time versus the response size for all the RDAP queries executed against the RIR's RDAP services. Pearson product-moment correlation coefficient was computed to assess the relationship between the response size and response time. A moderate positive correlation was found (Pearson r≈0.38, p-value≈0) which indicates that some larger responses took longer to be received.

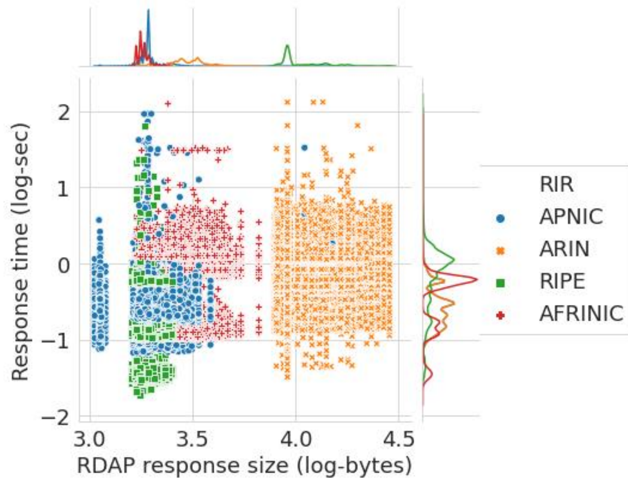

Figure 6.5 RDAP response size vs response time

7 TLD Registries' RDAP Response Time

Next, the response time for the RDAP services provided by various categorizations of TLD operators was analyzed. It is worth noting some relevant differences between the different servers. In particular, the response time was higher for RDAP servers providing country-level registration data (ccTLDs) than servers providing gTLD registration data. TLD registries were classified according to the type of TLD as stated in [IANA's root zone database](). As can be seen from Table 7.1, the average response time for ccTLDs was almost 1.4 seconds, while for gTLDs it did not reach even one second on average. This is due to the presence of heavy outliers within the measurements for ccTLDs.

Response time (sec)							
Type	**mean**	**std**	**min**	**50%**	**95%**	**99%**	**max**
ccTLD	1.36	1.33	0.10	1.18	2.81	3.17	208.21
gTLD	0.89	1.05	0.04	0.76	1.90	2.68	257.86

Table 7.1 Summary statistics of the RDAP response time depending on the TLD type

By looking at the boxplots in Figure 7.1, it becomes clear that RDAP servers handling ccTLD registration data were generally slower than those servers providing gTLD data.

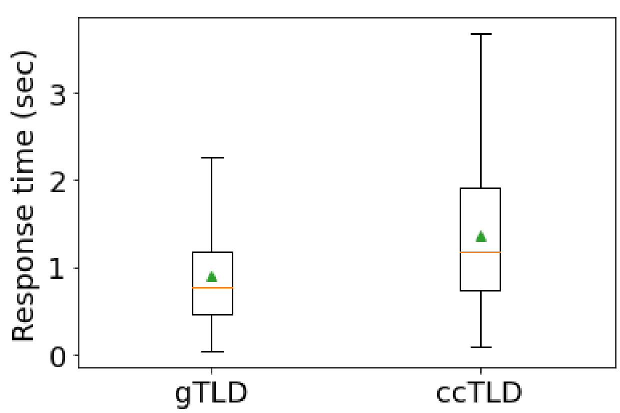

Figure 7.1. Response time distribution for RDAP servers depending on the TLD type (outliers have been removed from the boxplots)

7.1 Response Time Versus the Vantage Point

When analyzing the influence of the vantage point on the response time for the TLD's RDAP services, similar patterns to the previous case appeared. The queries run from the vantages points in Brazil, South Africa, and India are the ones that experienced the slowest RDAP response time (see Table 7.2).

	Response time (sec)						
location	**mean**	**std**	**min**	**50%**	**95%**	**99%**	**max**
BR	1.31	1.20	0.06	1.08	2.62	3.65	143.90
FI	0.86	0.85	0.10	0.75	2.15	2.82	156.91
FR	0.91	1.37	0.08	0.70	2.22	4.79	153.31
IN	1.21	0.76	0.22	0.95	2.66	3.58	74.17
JP	1.20	0.95	0.07	1.05	2.68	3.85	76.61
NL	0.84	2.04	0.04	0.61	2.13	4.61	257.86
ZA	1.32	0.97	0.30	1.13	2.59	3.52	148.42
SG	1.26	2.34	0.18	0.97	2.72	5.35	199.47
US1	0.77	0.54	0.04	0.66	1.81	2.47	16.70
US2	0.91	1.39	0.06	0.66	1.96	3.39	157.95

Table 7.2 TLD's RDAP response time depending on the location of the vantage point

7.2 Response Time Versus the Source IP Address Type

As aforementioned, a subset of the vantage points allowed us to perform RDAP queries using IPv6 source addresses. From these, the same queries were performed using IPv4 and IPv6 source addresses. As can be seen in Figure 7.3, no significant differences are found between the queries that used IPv4 compared to IPv6. Nevertheless, the response time for IPv4 queries was 20 milliseconds slower on average than queries effectuated over IPv6. Again, IPv4 queries suffered from more variance in the response time (std: 1.62 sec) than IPv6 queries (std: 0.58 sec)

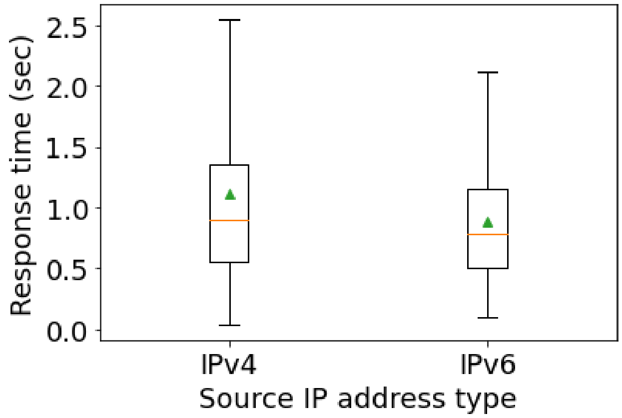

Figure 7.3 Response time per IP source address type (no outliers; triangles showing average response time)

7.3 Response Time Versus the Response Size

In terms of response size, gTLD's RDAP responses were almost double in size (~8 kbytes) than those provided by ccTLD's RDAP services (~4.3 kbytes). Again there was a huge variance among the different gTLDs (std: 5.3 kbytes).

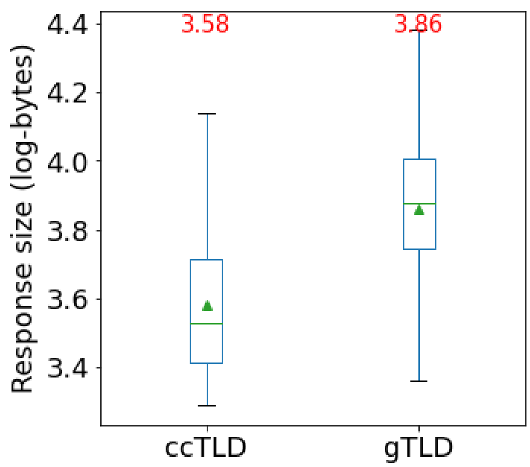

Figure 7.4 RDAP response size per TLD type (no outliers; triangles showing average response time)

To further investigate why some ccTLDs have higher latency than the gTLDs, the response size in bytes for a RDAP query is computed. Figure 7.5 shows these metrics against each other for each TLD type. Pearson's correlation coefficient (Pearson r=0.03, p-value=0.0) shows that there is no evidence that these delays in the RDAP responses are due to higher response sizes.

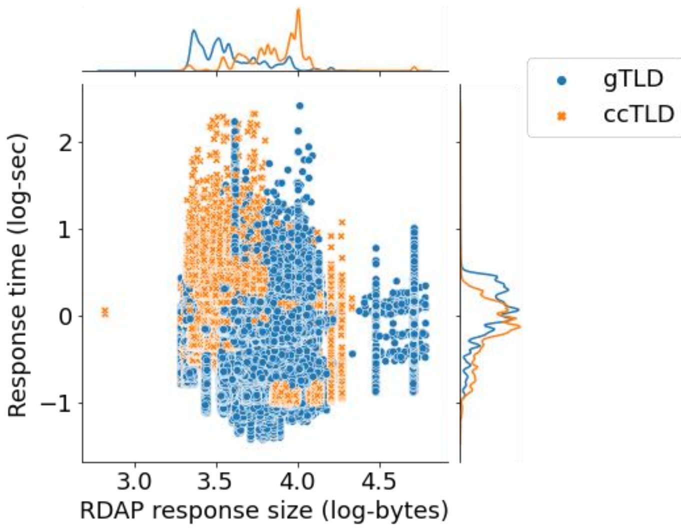

Figure 7.5 RDAP response size versus the response time

8 ICANN-Accredited Registrars' RDAP Response Time

While more than 2,400 ICANN-accredited registrars provide an RDAP service, only a total of 168 were probed. As described in the methodology section, these were registrars whose RDAP Base URL was present in IANA's bootstrap service and for which sample domains could be extracted from the .com zone file or from Spamhaus's passive DNS tool. Thus, registrars that had no domain in the .com zone file as of December 2020 were not included in the measurement.

The response times for the different registrars that belong to a particular family are grouped. A "registrar family" is a group of registrars that are under a common ownership or have a parent or subsidiary relationship organizational structure. Table 8.1 shows the statistics of the RDAP response time for the different registrar types. Registrars with family groups exhibit larger variance while the median response times are very similar.

	Response time (sec)						
Registrar Type	**mean**	**std**	**min**	**50%**	**95%**	**99%**	**max**
Family	1.10	3.06	0.03	0.75	3.02	8.08	940.40
Individual	0.98	1.48	0.03	0.79	2.29	5.36	168.88

Table 8.1 The RDAP response time, depending on whether a registrar belongs to a family of registrars

To further visualize the differences in response time among registrars' RDAP services, the top 50 registrars with the slowest response time were selected. Figure 8.1 shows the distribution of the response time for these registrars. As can be seen, the variance differs significantly among the different services.

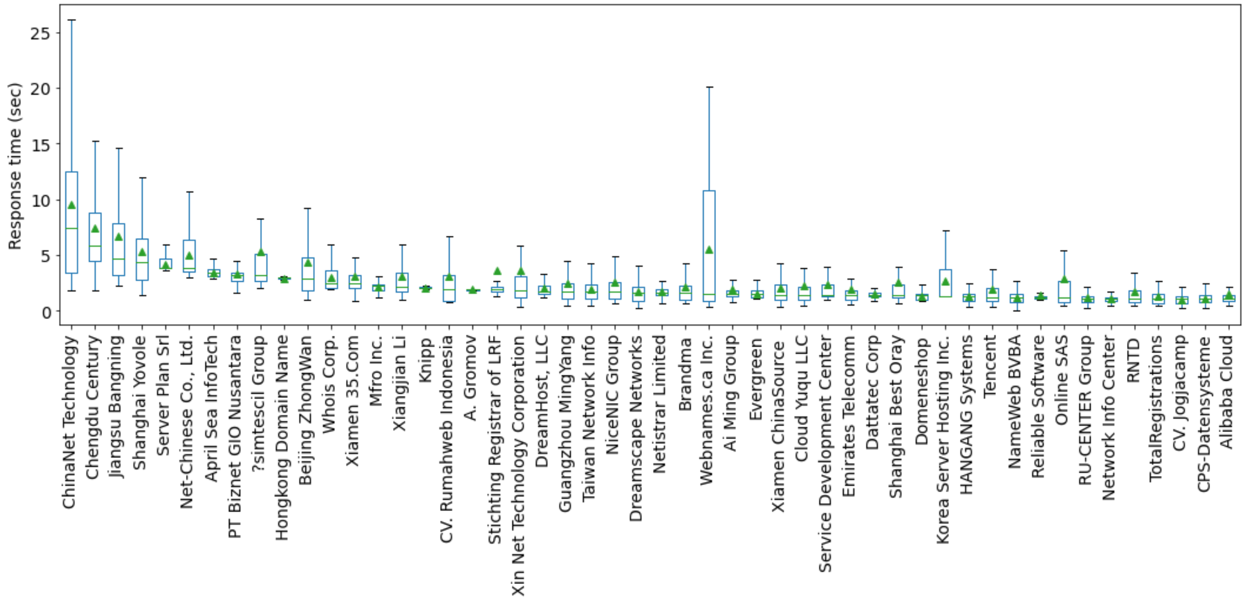

Figure 8.1 Query response time distribution for the top 50 slower registrar RDAP services (outliers have been removed from the boxplots)

8.1 Response Time Versus the Vantage Point

As in the previous cases, the vantage point from which the RDAP queries were executed affected the response time. Table 8.2 shows the summary statistics of the response time per vantage point.

	Response time (sec)						
Location	**mean**	**std**	**min**	**50%**	**95%**	**99%**	**max**
BR	1.59	2.37	0.07	1.10	3.62	8.19	128.98
FI	1.05	1.69	0.09	0.77	2.61	6.57	99.24
FR	0.96	1.72	0.06	0.68	2.43	6.11	132.22
IN	1.31	1.61	0.11	1.03	2.60	6.67	96.54
JP	1.31	1.51	0.07	1.11	2.51	6.46	91.85
NL	1.00	3.82	0.03	0.45	3.34	10.21	940.40
ZA	1.54	1.60	0.04	1.21	3.28	7.00	92.56
SG	1.32	2.13	0.05	1.03	2.71	7.74	160.76
US1	0.73	0.49	0.03	0.64	1.55	2.46	14.98
US2	0.96	1.66	0.04	0.65	2.30	6.12	165.67

Table 8.2 Registrar's RDAP response time depending on the location of the vantage point

8.2 Response Time Versus the Source IP Address Type

As can be seen in Figure 8.3, no significant differences were found between the queries that used IPv4 source addresses compared to IPv6 source addresses. Nevertheless, the response time for IPv4 queries was 17 milliseconds slower on average than queries effectuated over

IPv6. Again, IPv4 queries suffer from more variance in the response time (std: 2.89 sec) than IPv6 queries (std: 1.20 sec).

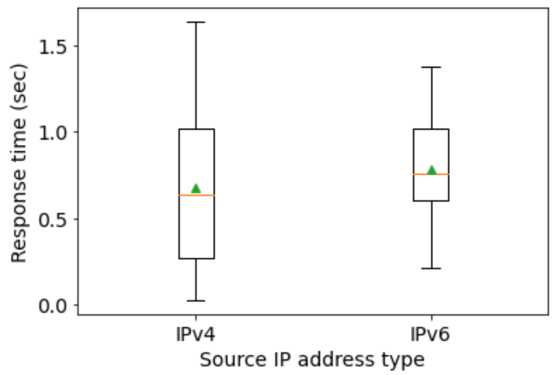

Figure 8.3 Response time per IP source address type (no outliers; triangles showing average response time)

8.3 Response Time Versus the Response Size

In terms of response size, the average RDAP response size of the queried registrars was around 4.5 kbytes. Again, there was a significant variance in response sizes among the different registrars (std: 2.3 kbytes). The fact that some registrars belong to a family did not make any difference in terms of the response size.

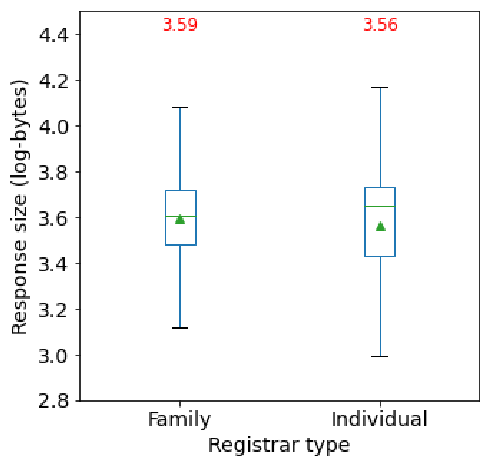

Figure 8.4 Response size per registrar family and individual registrar (outliers have been removed from the boxplots)

Finally, the relationship between response size and response time was investigated. Figure 8.5 shows these metrics against each other for each query and against a registrar's RDAP service. Pearson's correlation coefficient (Pearson r=0.005, p-value=0.0) shows that there is no relationship between response time and response size.

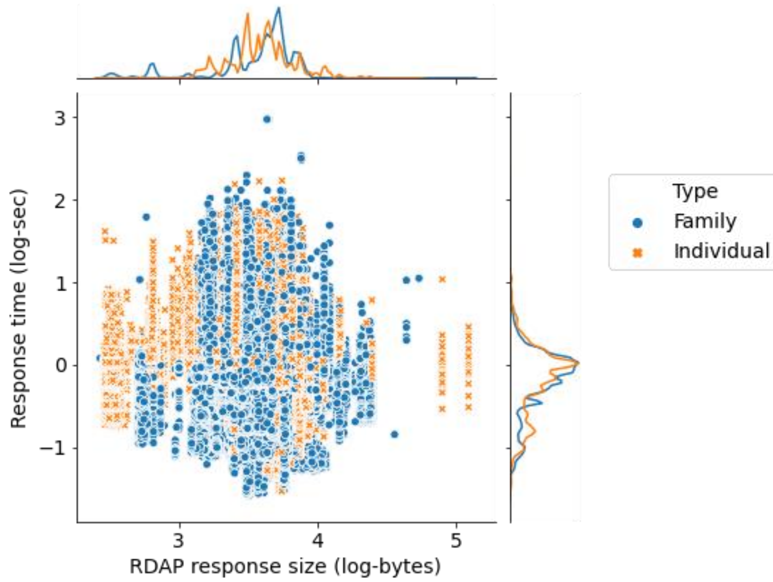

Figure 8.5 RDAP response size versus the response time

9 Limitations

- As with any active measurement study, this investigation suffers from the potential biases that a limited set of vantage points might create. While ten different vantage points distributed across multiple Autonomous Systems were used, it is unknown whether queries run from different Autonomous Systems could experience different response times.
- The study was purposefully designed to not stress the RDAP services. With a rate of 1 RDAP query every 5 minutes, the measurement should not have saturated any of the RDAP services. It remains unknown how the services would respond under heavier query loads. It is also possible that some services were under heavy use by other parties at the time our queries were made.
- For a minority of RDAP services, it was not possible to find a sample set of domains.
- Not all the vantage points allowed running queries over IPv6. The variance in response time might be impacted by the number of queries over IPv6.
- The validity of the RDAP responses was only assessed against the first level of conformance ("rdapConformance": ["rdap_level_0"]) and the presence of no errors.

10 Conclusions

The current deployment of RDAP services is diverse and, while 95% of all performed domain queries were answered under 4 seconds, some queries can take up to several minutes to be answered. Registrars' RDAP services are currently the slowest on average while RIRs' RDAP services are the fastest.

Several factors influenced the response time. The source location of the query had a significant impact on the response time. Queries originating from Europe and North America received faster responses than those from Asia or Africa. Queries executed over IPv6 had lower response times than those over IPv4. Response size did not seem to have a significant impact on the response time with the exception of ARIN's RDAP service whose responses were one order of magnitude larger than the other RIRs.

A Registrars' RDAP Errors

Based on the information available on IANA's bootstrap as of December 2020, multiple RDAP services did not reply to the queries. Table A.1 shows the RDAP Base URL of the registrars whose service did not respond to our queries.

RDAP base URL	Sample Query	Error
https://rdap.007names.net	https://rdap.007names.net/domain/purejoyparentingpopup.com	Connection refused
https://rdap.190.vip	https://rdap.190.vip/domain/sz-lexus.com	Internal server error
https://rdap.bb-online.com	https://rdap.bb-online.com/domain/lamarketingco.com	Connection timed out
https://rdap.bizcn.com	https://rdap.bizcn.com/domain/bjssnr.com	Connection timed out
https://rdap.conac.cn	https://rdap.conac.cn/domain/xn--fiq7vi3qcven98bizaw67l.xn--55qw42g	Connection timed out
https://rdap.dns.business	https://rdap.dns.business/rdap/domain/cfm-za.com	Connection refused
https://rdap.domainthenet.com	https://rdap.domainthenet.com/domain/alliberthome.com	Not found
https://rdap.nicproxy.com	https://rdap.nicproxy.com/domain/53crew.com	Runtime Error
https://rdap.pavietnam.vn	https://rdap.pavietnam.vn/domain/quynhon.review	Bad request
https://rdap.safenames.legal	https://rdap.safenames.legal/domain/goodsbusiness.com	Connection timed out
https://rdap.sfn.cn	https://rdap.sfn.cn/rdap/domain/oiasmart.com	Internal server error
https://rdap.smallworldregistrar.com	https://rdap.smallworldregistrar.com/domain/hashtaggethired.com	sslv3 alert handshake failure
https://rdap.west.cn	https://rdap.west.cn/domain/beer.beer	Bad request
https://rdap.zxnic.cn	https://rdap.zxnic.cn/rdap/domain/motuan.vip	Unsupported media type
https://whois.dns.com.cn	https://whois.dns.com.cn/rdap/domain/FEISHENGBO.COM	Connection timed out
https://whois.xz.com	https://whois.xz.com/rdap/domain/bananamatrix.com	404 Error
https://www.yesnic.com	https://www.yesnic.com/rdap/domain/clab.pro	Empty reply

Table A.1 Errors obtained while querying RDAP's registrar service