Examining the Centralization of Email Industry: A Landscape Analysis for IPv4 and IPv6

Luciano Zembruzki*, Arthur Selle Jacobs*, Ricardo Jose Pfitscher†, Lisandro Zambenedetti Granville*

*Institute of Informatics (INF) - Federal University of Rio Grande do Sul (UFRGS)
Porto Alegre, Brazil
{lzembruzki, asjacobs, granville}@inf.ufrgs.br

†Department of Mobility Engineering - Federal University of Santa Catarina (UFSC)
Joinville, Brazil
ricardo.pfitscher@ufsc.br

***Abstract*— Centralization of key Internet services, including email, can result in privacy and security concerns and increase the number of single points of failure. This paper measures and analyzes a large-scale dataset of email providers gathered from MX records of top-level domains. The findings reveal the concentration of email infrastructure providers for each TLD and identify the most significant providers in the market. The paper also demonstrates that the IPv6 adoption increased the centralization of email servers. The research contributes to the state-of-the-art by thoroughly examining email infrastructure centralization and identifying potential areas for future research.**

I. Introduction

In recent years, the centralization of essential Internet services such as the Domain Name System (DNS), Web hosting, and email has become more prevalent. While centralization can lead to economic benefits by sharing IT infrastructure costs among customers, concerns about its impact on privacy and security have been raised by academics, operators, non-profit organizations, standardization bodies, and policymakers [1]–[11]. Specifically, the transfer of clients' services and data to third-party servers has led to concerns about privacy and security [4] [5]. Moreover, centralization increases dependencies on a single provider (single point of failure) [9]–[11] and can lead to catastrophic failures in the event of service disruptions. Failures at this single point can have a ripple effect, impacting the availability of services for all organizations that share the same infrastructure [12] [13]. Given the growing reliance on centralized Internet services, examining the potential risks and drawbacks associated with this trend is crucial.

Two notable examples of the risks associated with centralization are the Dyn cyberattack in 2016 [14], and the Google Mail Server Outage in 2020 [15] [16]. The Dyn attack targeted a DNS provider through a botnet of Internet of Things (IoT) devices, resulting in widespread disruption of Internet services and major websites becoming inaccessible for several hours. In the case of the Google Mail Server Outage, a hardware issue in Google's data center affected millions of users worldwide for several hours [15] [16]. These examples illustrate the potential impact of centralization on the availability of services and the need for organizations to carefully evaluate the security measures and policies of third-party providers. To ensure that their data is protected and compliant with relevant laws and regulations, organizations should consider implementing measures such as regular security audits, encryption, and data backup strategies [1] [3].

As concerns about Internet centralization have grown, several research initiatives have investigated various aspects of the problem. For instance, some researchers have focused on service centralization and investigated how Web hosting providers operate in today's Internet market [17]. Other researchers have examined the centralization of the Internet Domain Name System (DNS) industry to understand how DNS providers operate across the Internet and whether they contribute to further centralization [18] [7]. Additionally, in terms of email service centralization, Liu *et al.* [19] developed a methodology for mapping domains to mail service providers and utilized this approach to documenting the dominance of email service providers and their growing role over the years.

In this paper, we measure and analyze a large-scale dataset of email providers to thoroughly examine the concentration of email service infrastructure. To accomplish this, we use a methodology that gathers data from Mail Exchanger - MX records of top-level domains (TLDs) and analyzes email infrastructure provider distribution. Our dataset comprises MX records collected and enriched with additional information for a list of domains from the Tranco List [20] as of 2023. Our findings reveal the concentration of email infrastructure providers for each TLD and identify the most significant providers in the market. Furthermore, we demonstrate that the IPv6 adoption has increased the centralization of email servers, with top servers' geolocation migrating from America to Europe.

The remainder of this paper is organized as follows. Section II discusses related work on measuring various Internet centralization and consolidation aspects and positions our work within the state-of-the-art. Section III describes our analysis methodology, limitations, and datasets. Section IV presents our results, analyzing the concentration of email infrastructure providers for each TLD and demonstrating how IPv6 adoption has affected the centralization of email servers. Finally, in Section V, we summarize our findings and outline future research directions.

II. Related Work

Several studies have demonstrated the centralization trend in critical Internet services, including *Web Hosting*, *DNS*,

979-8-3503-0048-2/23/$31.00 ©2023 IEEE

and *email*. Zembruzki *et al.* [17] conducted a comprehensive analysis of the Web hosting industry and found that five large US hosting providers controlled over a third of 150 million evaluated domains from 19 TLDs. The study also revealed an increasing centralization trend between 2016 and 2021, with Google emerging as the largest provider in 2021, hosting 18% of all domains. The authors also discovered that geographic proximity and shared language ties significantly influenced the hosting sector. The same authors conducted another analysis, now over the DNS industry [18], and revealed that the Top-5 DNS providers account for over 20% of all domains. In contrast, the Top-100 providers account for nearly 80% of the examined IPv4 domain namespace.

Radu *et al.* [21] investigated the growing consolidation trends in the recursive DNS services industry. They analyzed the industry's evolution over the past decade and provided empirical evidence of changes between 2016 and mid-2019. The research relied on active measurements from 100,000 mobile devices and probes and revealed that public DNS resolvers handled over 50% of the total DNS requests in the first half of 2019. The authors also found that Google and Cloudflare dominate the market, managing half the market share. Additionally, they noted that the future of DNS in browsers looks similar, with Google strengthening its position and limiting competition opportunities.

Kashaf *et al.* [8] analyzed the presence and impact of third-party dependencies in three infrastructure services: DNS, CDNs, and certificate revocation checks by CAs. The methodology considered both direct and indirect dependencies. The findings showed that 89% of Alexa's Top 100,000 websites rely heavily on third-party DNS, CDN, or CA providers. Therefore, these websites would experience service disruptions in case of failures in these third-party providers. The study also revealed that third-party service use is highly concentrated, with the leading Top-3 providers of CDN, DNS, or CA services affecting between 50-70% of the Top 100,000 websites. However, their analysis of DNS dependencies relied solely on NS record labels (*e.g.*, ns1.example.com), which can mask cases of centralization as a single IP address can host multiple name servers.

In terms of email services, Liu et al. [19] conducted a study on the centralization of email providers used by organizations. They found that most organizations rely on third-party providers for email services, concentrating on a small number of large providers. This centralization increases the risk of single points of failure and data privacy concerns. The authors emphasized the importance of carefully evaluating third-party providers' security measures and policies to protect data and comply with relevant laws and regulations. The results of their study highlight the need for a more diverse and distributed email infrastructure, which would mitigate the risks associated with centralization.

In this study, we contribute to the state-of-the-art by examining the centralization of email infrastructure in organizations using a dataset of email servers from various top-level domains. We analyze the concentration of power among email providers and identify potential risks and challenges associated with centralization. Furthermore, we investigate the impact of IPv6 adoption on centralization, which has yet to receive much attention in the literature. In summary, our study differentiates from the related work in two aspects: we focused on the infrastructure hosting the email services. We provided a comprehensive analysis of the impact of IPv6 adoption on centralization.

III. Methodological Overview

In this section, we provide a detailed explanation of the methodology and data sources used to analyze the current state of the email industry. We also describe our assumptions and limitations.

A. Methodology

Due to the need for legal agreements with registry operators, access to zone files containing the entire global DNS namespace is generally limited, making it impractical to cover the entire namespace, particularly when considering certain ccTLDs. Thus, to conduct a large-scale analysis, we rely on a DNS measurement using rather the email records of domains from the Tranco list [20]. The Tranco list, which ranks the most visited domains worldwide, provides a representative sample of the DNS namespace and is updated daily. The list includes domains from various gTLDs such as `.com`, `.net`, and `.org`, as well as several ccTLDs for many continents (such as `.br`, `.uk`, and `.nl`).

Figure 1 illustrates the leveraged email measurement method. Initially, the MX records for each domain in the Tranco list are queried to identify the email servers associated with them. Then, the IPv4 and IPv6 addresses linked with each MX record for each domain are searched, as specified in RFC 1035. The A and AAAA records, respectively, as per RFC 1034 and RFC 3596, are used to find the IPv4 and IPv6 addresses. For instance, when analyzing the domain `ufrgs.br`, the MX record is `mx1.ufrgs.br`. The corresponding A record for the IPv4 address is `143.54.1.201`, and the AAAA record for the IPv6 address is `2804:1f20:0:1000::201`.

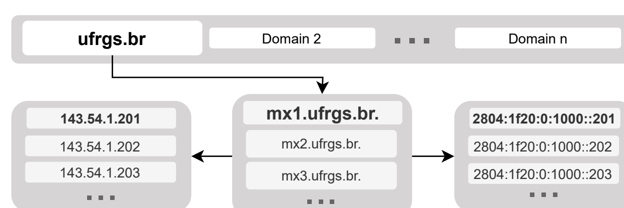

Figure 1: Collected data

After retrieving the MX records and corresponding IP addresses, we enrich the data using the dataset schema depicted in Figure 2. To begin with, we incorporate information about the Autonomous Systems (ASes) that announce the IP addresses, utilizing data from the CAIDA prefix-to-AS mapping [22]. We also map the AS numbers to organizations and countries based on the AS-to-organization dataset [23]. Furthermore, we associate each country code with one or more official languages, drawing on a country and language dataset [24]. Lastly, we use the database from [25] to identify the geographic location of each email server.

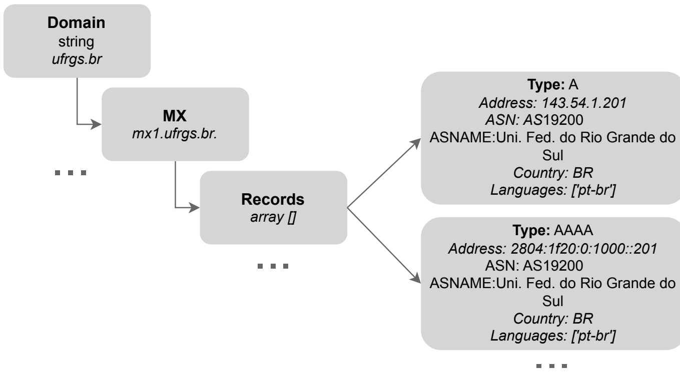

Figure 2: Dataset Schema

Enriching the collected data allows for a more comprehensive analysis of the DNS landscape by providing information on the organizations and countries responsible for hosting email servers associated with top domains. Additionally, collecting IPv4 and IPv6 addresses associated with each domain enables an analysis of the impact of IPv6 adoption on centralization. Furthermore, the geolocation information and country map enable an investigation of how organizations use services hosted in foreign countries.

B. Datasets

We collected the data from the Tranco list and run the collection methodology in January 2023. Table I summarizes the dataset when considering all aggregated queries.

Year	Domains	with MX	MX(v4)	MX(v6)	AS(v4)	AS(v6)
2023	1M	481k	409k	113k	15k	1.8k

Table I: Aggregated dataset, measured on January of 2023.

The dataset from the Tranco list [20] comprises 1 million domains, of which 481k have MX records (48.1%). These MX records contain 409k entries that resolve to IPv4 addresses (MX(v4)) and 113k entries for IPv6 addresses (MX(v6)). The columns AS(v4) and AS(v6) indicate the number of Autonomous Systems that announce the IPv4 and IPv6 addresses, respectively. The IPv4 addresses are announced by 15k ASes, while the IPv6 addresses are announced by 1.8k ASes. For replication purposes, we published our dataset on GitHub[1].

C. Assumptions and Limitations

Our methodology includes assumptions and limitations:

Email hosting: For simplicity, we assume that an MX record for a domain implies the existence of the email server. This assumption will break for domain names that present the record but do not have an email server listening on the relative address.

[1]https://internet-centralization.github.io/

Email popularity: Although the Tranco list contains the most visited domains worldwide, it cannot imply that these domains are the most used email services. However, given the access restrictions to ISP data and the lack of public data on SMTP traffic, we assume that the domains from the Tranco list are representative enough for our study.

Parket domains: If a domain name represents a non-responsive or parking domain, we still consider the infrastructure if the domain presents MX records - the classification of this typology of domains is left for future work.

Server geolocation: The accuracy of server geolocation varies, and several factors can affect it. However, despite the varying geolocation accuracy, server geolocation can still provide valuable insights for trend and pattern analysis.

Single vantage point: Even if an email provider operates address spaces in multiple countries, the AS number of its headquarters country remains the same. We collect our measurements from a single vantage point in Brazil, which can introduce a bias towards "nearby" MX records. This bias could increase the concentration of IP address space in America.

IV. RESULTS

The results presented in this section aim to uncover the extent of centralization in the email industry with regard to the use of IPv4 and IPv6 addresses. First, we show the percentage of domains with email servers hosted by the top email providers' Autonomous Systems (ASes). Next, we investigate the concentration of top providers per Top-Level Domain (TLD). Similarly, we examine the market share, first for the aggregated dataset and then for each TLD. We also investigate the location profile of the dominant companies in each TLD. Finally, we present a geolocation mapping of the servers.

A. Top Email Providers

To determine if there is significant concentration of domains within a particular group of ASes, we calculated the percentage of domains announced by the top 1, 2, 3, 4, 5, 10, 15, 20, 25, and Top 100 ASes of email providers. Figure 3 shows the percentage of domains hosted by the top ASes for both IPv4 and IPv6 addresses.

The results presented in Figure 3a indicate a significant degree of centralization in the email infrastructure, specifically for domains that resolve MX records to IPv4 addresses. Approximately 50% of domains from the Tranco list are associated with only ten ASes belonging to email providers. Upon closer examination, the results indicate that a single email provider AS (top 1) accounts for 20.6% of all domains with email servers, while the top 5 email provider ASes account for 41.6% of all domains. Moreover, the data shows that only 100 providers concentrate 70% of all email domains from the Tranco list, indicating a high level of centralization in the email industry.

The results presented in Figure 3b reveal significant centralization in the email infrastructure for domains that resolve MX records to IPv6 addresses, despite being less representative than IPv4 addresses (23.4% of total MX records). The top email provider AS controls over 72.5% of the domains in our dataset, and the top five account for 87% of the domains, indicating a high level of concentration. These results are even more striking than those for IPv4 centralization. While in IPv4, the top 100 email providers host 70% of domains, in IPv6, only ten email provider ASes account for more than 90% of all domains. This concentration suggests that a few organizations have significant dominance in the email infrastructure for IPv6.

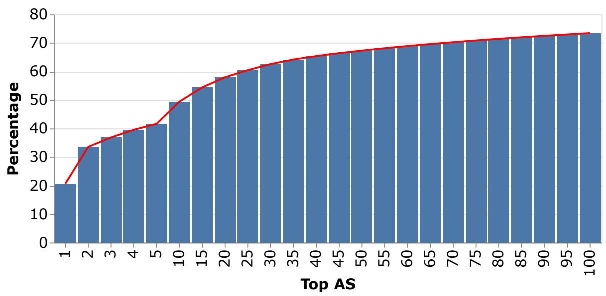

(a) IPv4

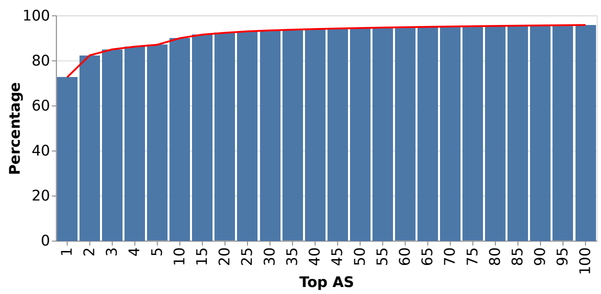

(b) IPv6

Figure 3: Top MX Providers Concentration

B. Top Email Providers per TLD

Next, we explore the email infrastructure providers for each TLD to investigate whether the concentration of email providers is similar or different across TLDs. To accomplish this, we calculate the concentration of the top 5, 10, 20, and 100 email provider ASes, similar to our previous analysis.

Figure 4 reveals that the concentration of top MX infrastructure providers varies according to TLDs and IP versions. We observe non-uniform concentration among TLDs in the case of IPv4 (Figure 4a). While the top 5 email provider ASes account for 43.47% of `.br` domains and 64.97% of `.cn` domains, more than 100 providers hosts 80% of the `.com` and `.net` gTLDs. When we expand our view to the top 10 email provider ASes, we find that these providers maintain more than 50% of `.uk`, `.nl`, and `.jp` domains, indicating that the top 10 providers have a firm grip on the MX infrastructure for these TLDs. With respect to IPv6 (Figure 4b), we note that most TLDs have approximately 85% to 90% domains concentrated in just five providers. We highlight the `.nl`, `.de`, and `.net` TLDs as having the lowest indexes, but still exhibiting impressive levels of concentration for IPv6.

C. Email Providers Market Share

Next, we analyze the market share of email hosting for each individual AS across our aggregated dataset. Figure 7 displays a

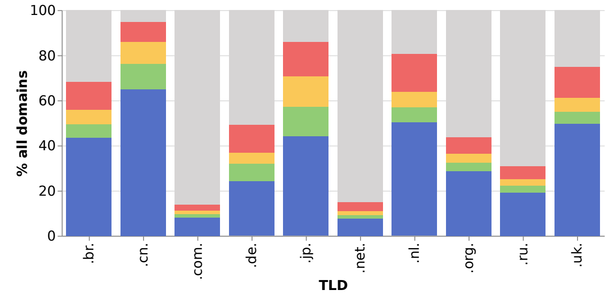

(a) IPv4

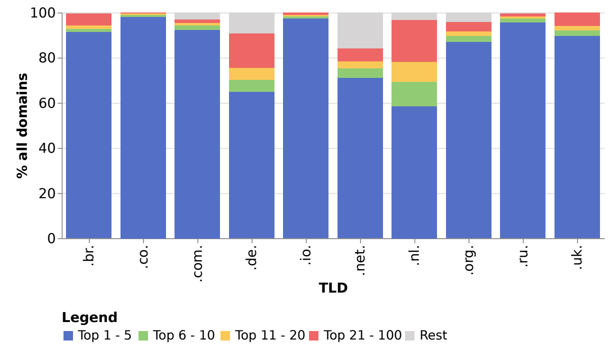

(b) IPv6

Figure 4: Top MX Providers Concentration per TLD

chart on the right indicating the percentage of domains for each company and a chart on the left showing the absolute number of domains for each company. Figure 3a depicts the email hosting market for IPv4 addresses, with Google, Microsoft, and Amazon dominating the market. Google holds the lead with a 20.65% share, closely followed by Microsoft with 12.89%. Amazon is in third place, with a market share of 3.31%. Concerning the absolute number of hosted domains, Google has 84,474 domains, Microsoft has 52,755 domains, and Amazon has 13,529 hosted domains. The remaining companies hold a smaller market share.

When examining the email hosting market based on IPv6 data (Figure 3b), it is apparent that Google holds an even more dominant position, having 72.56% of the market. This represents a significant increase compared to the IPv4 email hosting data, where Google had a 20.65% share. Except for the first position, the rankings of the remaining companies, however, were largely different between IPv4 and IPv6. Yandex Oy, for instance, rose from fourth to second, with a 9.70% market share. This suggests that some companies may be more advanced in implementing IPv6, even outpacing industry giants like Microsoft and Amazon.

D. Email Providers Market Share per TLD

In Figure 6, we can see the market share per TLD in both IP versions. Figure 6a shows the IPv4 data, where we observe that Google and Microsoft dominate the market by frequently alternating in the first and second positions in most TLDs, with the exception of `.ru` and `.cn`, where national companies from Russia and China, respectively, hold the majority of the market

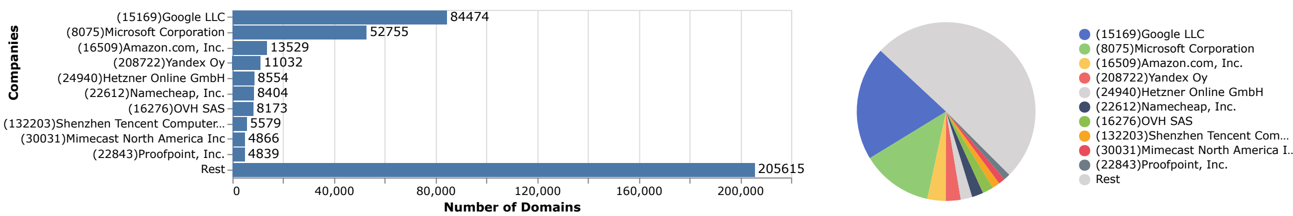

(a) IPv4

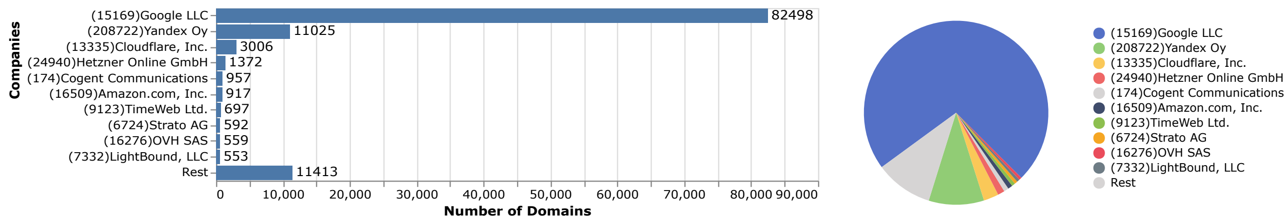

(b) IPv6

Figure 5: Market Share of Email Providers on aggregated dataset

share. This behavior may be explained by political issues. It is worth noting that the `.net`, `.ru`, and `.de` TLDs have around 50% of the market share distributed among smaller companies, which we label as *Rest* in the charts.

The market share for IPv6 (Figure 6b) confirms our previous findings of centralization, with one or two companies dominating most TLDs, except for `.nl` and `.de`, where minor companies have a significant presence. Google continues to monopolize the market in most TLDs, with market shares exceeding 75%, and even reaching up to 90% in the `.io` and `.co` TLDs. It is also interesting to highlight the `.ru` TLD, where Yandex, a Russian company, holds around 75% of the market share. Yandex also appears in the second position in other TLDs, but with less representation.

E. Top Email Providers per Characteristics

Next, we analyze the geographic distribution of email companies by looking up the company's origin for each AS used for email hosting. We map the domain names associated with each AS to their corresponding countries. It is important to note that this process differs from geolocating hosting IP addresses, as explained in Section III.

For each top-level domain (TLD), we tally the number of domains per country and classify each country based on four priority categories. These categories are ranked by order of precedence: *local* (if the country matches the TLD), US-based (given the prevalence of large cloud companies in the US), same official language, and *rest* (all other domains). Each AS/country is labeled based on a single category, even if multiple categories apply. Note that gTLDs are not labeled as "local" and "language" because they are global. The analysis of IPv4 results (Figure 7a) shows that the `.cn` TLD has the highest percentage of domains relying on local email providers, with 65.69%, followed by `.jp` with 59.21%, and `.ru` with 55.53%. Conversely, the `.br` TLD has the highest percentage of US-based email providers with 59.29%, followed by `.uk` with 59.34% and `.de` with 27.23%. The gTLDs also have a high concentration of US providers as well.

There is no substantial change with IPv6 (Figure 7b), as US-based companies continue to dominate most TLDs. However, there are a few exceptions, such as `.de` and `.nl`, where there is a higher proportion of local companies. Notably, the TLDs `.co`, `.com`, and `.io` have a strikingly high percentage of domains relying on US-based companies, with 93% of domains falling in this category.

F. Top Email Providers Geo-location

Finally, for an in-depth understanding of the physical distribution of email infrastructure, we analyze the geographic locations of email servers. This analysis can help identify patterns of centralization or decentralization of email servers and provide insights into factors that influence geographical distribution. Figure 8 presents a global map of the geolocation of email servers. For clarity, we plotted only the top 10 email providers and grouped the remaining servers as "Rest". The size of the circles represents the number of domains each server hosts.

Analysis of the geographic distribution of IPv4 email providers (Figure 8a) reveals a concentration of servers in Europe and North America. Google dominates the market with a significant concentration of domains allocated to servers in the United States. In contrast, Yandex concentrates substantial number of email domains on servers in Russia. The domains hosted by Amazon, however, are more evenly distributed between the United States and Europe.

Looking at the geographic distribution of IPv6 email providers (Figure 8b), we observe a continued trend towards centralization, similar to what we saw with IPv4. Interestingly, we also observe a shift in the concentration on Google servers from the US to Europe. Such behavior could indicate a greater adoption of IPv6 in Europe or be related to internal decisions of the company. Yandex, on the other hand, maintains its servers

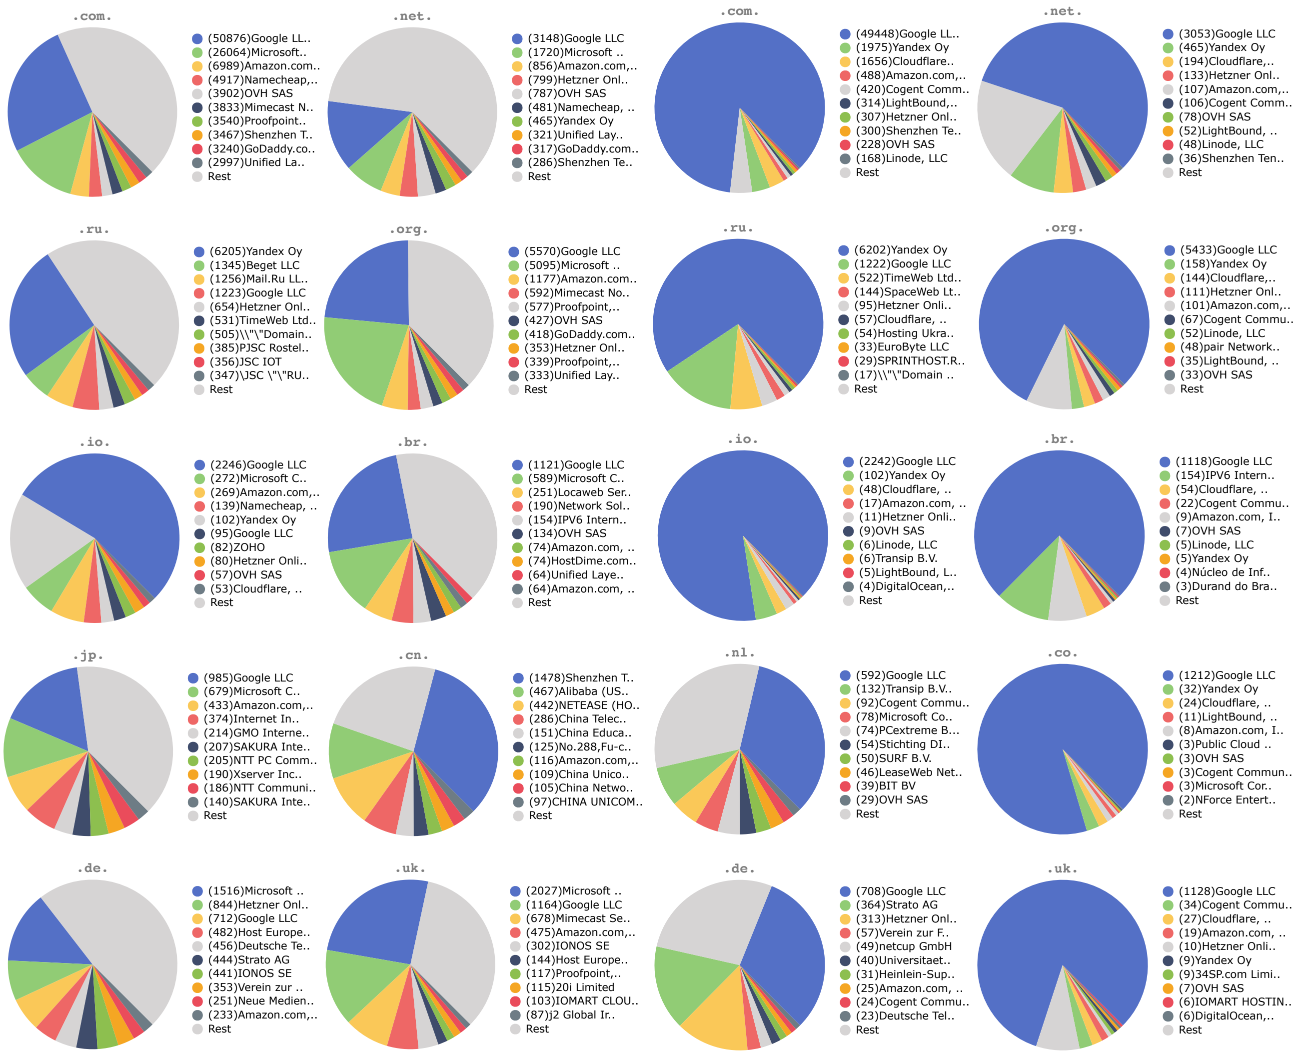

(a) IPv4 (b) IPv6

Figure 6: Market share of email providers per TLDs

in Russia for both IPv4 and IPv6. Finally, it is worth noting that Google has a significant presence in South America, with servers centralized in Chile.

V. Conclusions and Future Work

In recent years, there has been a growing concern among various sectors of society regarding the increasing consolidation and centralization of the Internet. The concentration of many essential Internet resources among a small number of providers is one of the critical concerns raised by this trend toward consolidation of *infrastructure*, *traffic*, and *services*.

In this paper, we investigated the extent of infrastructure centralization in the email industry using the Tranco list as a reference for domains. We measured the MX records for them, mapping each MX AS to its parent organization and estimating each email provider's market share. Our findings revealed that only some companies dominate the industry, with the Top 1 to 5 email providers accounting for more than 40% of all domains, and the Top 1 to 100 providers hosting around 70% of the total. The cases where MX records resolve to IPv6 addresses are even more concerning: the top 5 providers account for 87% of domains. Also, our research showed that Google dominates the global market share, holding 25% of the domains that use IPv4 and around 75% of the domains under IPv6. Finally, the geolocation study raised another concern: most domains run on top of servers physically located in the same region, which may imply outages in the case of failures.

As future work, we aim to continue our analysis of the email industry's infrastructure centralization by establishing a representative ranking for email domains based on SMTP, POP3, and IMAP traffic. While we used the Tranco list as a reference for domains, it is primarily based on DNS and Web requests. Therefore, a more representative ranking that includes email traffic data is necessary to provide a complementary view of the email industry's centralization. To achieve this, we plan to rely on data sources from multiple vantage points.

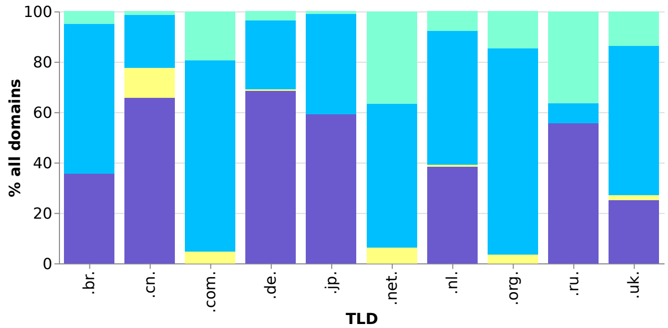

(a) IPv4

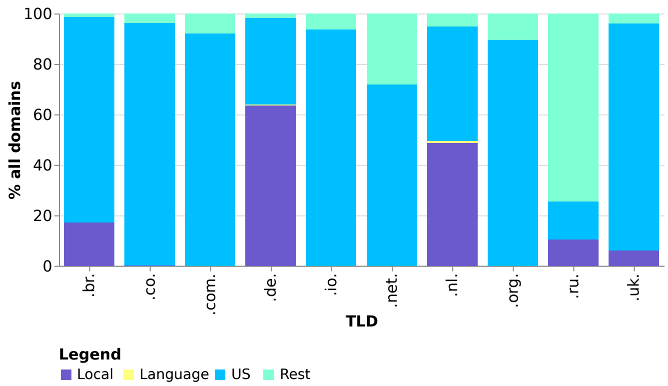

(b) IPv6

Figure 7: Top Email Providers Concentration per Characteristics

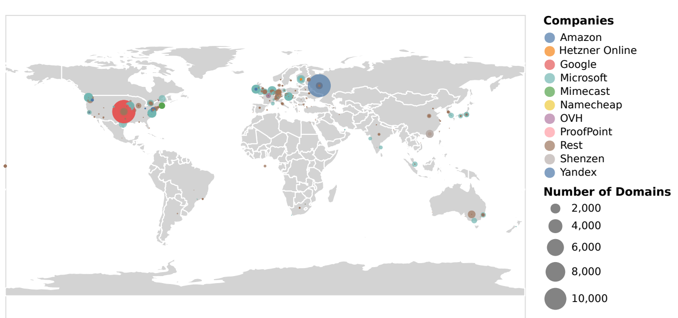

(a) IPv4

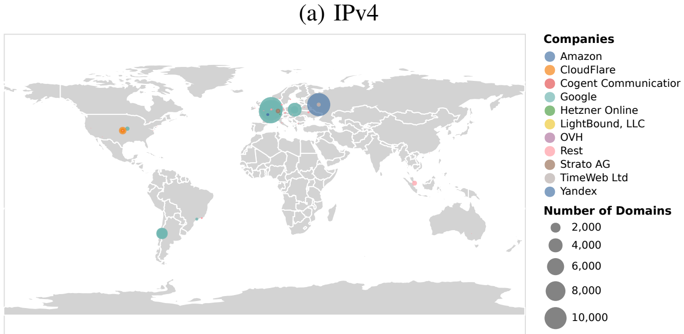

(b) IPv6

Figure 8: Top Email Providers Geo-location

Acknowledgements

This work was supported in part by CNPq procs. 423275/2016-0 (Universal), 316662/2021-6 (PQ), and 88887.480774/2020-00, by FAPESP procs. 2020/05152-7 (PROFISSA), 2015/24494-8 (BigCloud), by FAPERGS procs. 16/2551-0000488-9 (Green Cloud), and by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

References

[1] Jari Arkko. The influence of Internet architecture on centralised versus distributed Internet services. *Journal of Cyber Policy*, 5(1):30–45, 2020.
[2] Jari Arkko. Centralised Architectures in Internet Infrastructure. Internet Draft, November 2019.
[3] Jari Arkko et al. Considerations on Internet Consolidation and the Internet Architecture. Internet Draft, July 2019.
[4] Bruce Schneier. Censorship in the Age of Large Cloud Providers, 2018.
[5] Internet Society. Consolidation in the internet economy. Internet Society https://future.internetsociety.org/2019/, 2019.
[6] Cecilia Kang and David McCabe. Lawmakers, united in their ire, lash out at big tech's leaders. *New York Times*, July. 29 2020.
[7] Giovane C. M. Moura, Sebastian Castro, Wes Hardaker, Maarten Wullink, and Cristian Hesselman. Clouding up the internet: How centralized is dns traffic becoming? In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, page 42–49, New York, NY, USA, 2020. Association for Computing Machinery.
[8] Aqsa Kashaf, Vyas Sekar, and Yuvraj Agarwal. Analyzing Third Party Service Dependencies in Modern Web Services: Have We Learned from the Mirai-Dyn Incident? In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, page 634–647, New York, NY, USA, 2020. Association for Computing Machinery.
[9] Shoshana Zuboff. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile books, 2019.
[10] Tim Wu. *The Curse of Bigness: Antitrust in the New Gilded Age*. Columbia Global Reports, 2018.
[11] Steven Erlanger et al. Europe feels squeeze as tech competition heats up between u.s. and china. *New York Times*, Sep. 11 2020.
[12] Synthia Wang, Kyle MacMillan, Brennan Schaffner, Nick Feamster, and Marshini Chetty. A first look at the consolidation of dns and web hosting providers. *arXiv preprint arXiv:2110.15345*, 2021.
[13] Trinh Viet Doan, Justus Fries, and Vaibhav Bajpai. Evaluating public dns services in the wake of increasing centralization of dns. In *2021 IFIP Networking Conference (IFIP Networking)*, pages 1–9, 2021.
[14] DynDNS. https://dyn.com/dns/, 2018.
[15] Rishi Iyengar. Gmail suffers another outage | cnn business, Dec 2020.
[16] Alex Hern. Google suffers global outage with gmail, youtube and majority of services affected, Dec 2020.
[17] Luciano Zembruzki, Raffaele Somesse, Lisandro Zambenedetti Granville, Arthur Selle Jacobs, Mattijs Jonker, and Giovane Moura. Hosting industry centralization and consolidation. In *IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2022.
[18] Luciano Zembruzki, Arthur Selle Jacobs, and Lisandro Zambenedetti Granville. On the consolidation of the internet domain name system. In *GLOBECOM 2022 - 2022 IEEE Global Communications Conference*, pages 2122–2127, 2022.
[19] Enze Liu, Gautam Akiwate, Mattijs Jonker, Ariana Mirian, Stefan Savage, and Geoffrey M. Voelker. Who's got your mail? characterizing mail service provider usage. In *Proceedings of the 21st ACM Internet Measurement Conference*, page 122–136, New York, NY, USA, 2021. Association for Computing Machinery.
[20] Tranco List. Tranco. Tranco List https://tranco-list.eu/, 2023.
[21] Roxana Radu and Michael Hausding. Consolidation in the dns resolver market – how much, how dangerous? *Journal of Cyber Policy*, 2020.
[22] CAIDA. Routeviews Prefix to AS Dataset for IPv4 and IPv6, 2020.
[23] CAIDA. Inferred AS to Organization Mapping Dataset. https://www.caida.org/data/as-organizations/, 2017.
[24] Annexare Studio. Countries, languages & continents data. https://github.com/annexare/Countries.
[25] MaxMind. Geoip2 databases. MaxMind https://www.maxmind.com/en/geoip2-databases, 2023.