From Diversity to Dominance: The Consolidation and Centralization of Email Hosting in European ccTLDs

Tobias Seijsener
tseijsener@os3.nl
Research Project 2 - Security and Network Engineering - University of Amsterdam

Supervisors: Willem Toorop and Koen van Hove (*NLnet Labs)*

***Abstract*—This paper examines the consolidation and centralization of email hosting in 5 European Country Code Top-level Domains (ccTLDs). We use OpenINTEL's forward Domain Name System (DNS) measurements to inspect the Mail Exchanger (MX) records within these zones, normalize this data and discover centralization trends. We see a growing number of domains utilizing the same hosters, creating scenarios where the top 5 providers account for up to 70.1% of the email hosting across a single ccTLD. We highlight a trend of centralization among all analyzed zones and a declining share of smaller providers reducing the general diversity of email hosting providers. We also find that hosting diversity is influenced by providers from other countries that share the same language, as well as the geographical distance to the zone's country.**

I. Introduction

The consolidation and centralization of the Internet have long been a debated topic. Researchers [1], policymakers [2], standardization bodies and others alike are worried about a small set of organizations controlling bigger and bigger parts of the internet, with email being a striking example.

Email was originally designed as a decentralized way for organizations, people and entities to exchange mail with each other over a network of networks – The Internet. The underlying protocols for today's email were introduced in the fall of 1981, making them over 40 years old [3]. Today, email is still omnipresent. However, there is one big difference. GMail, Outlook, iCloud Mail, Yandex Mail, and other brand names have become synonymous with the term "email". Society is increasingly moving to these hosted digital infrastructures, which are largely managed by a small number of parties [4].

Multiple studies have been conducted on Internet centralization, with researchers analyzing the centralization within web hosting [5] and observing a heavily centralized market. A similar phenomenon was observed in a study regarding DNS centralization [6] and in an additional study regarding DNS centralization for the `.nz` and `.nl` zones [7].

In the current landscape of internet centralization research, no work has observed whole DNS zones, using historical data to discover centralization trends in the email hosting market, with a specific focus on European ccTLDs. Our research uses the MX records of domains in a top-level domain to figure out which providers are being used and uses this information to spot trends among the historical data points. This paper aims to identify these trends, quantify the extent of centralization, and examine the role of different hosting providers in this process.

In section II we outline the used terminology and technologies, in section III we dive into the previous research on the topic of internet centralization. Following that, we outline our used methodology in section IV, where we in addition make assumptions and explain our dataset. Following this we will share the results in section V, followed by section VI where we highlight trends and commonalities and discuss them. To conclude the paper we will share our takeaways in section VII and share our reflections for future research in section VIII.

The research leading to these results was made possible by OpenINTEL[1], a joint project of the University of Twente, SIDN, NLnet Labs and SURF.

A. Research Questions

Research question: *How has the centralization of email hosting evolved across different European ccTLDs, and what are the implications for digital sovereignty?*

- **Subquestion 1:** *What historical trends in email hosting centralization can be identified through DNS record analysis within these ccTLDs?*
- **Subquestion 2:** *How does email hosting centralization affect the control of European nations over their digital communication infrastructure?*
- **Subquestion 3:** *How does the level of centralization differ between these ccTLDs?*

II. Background

We first describe consolidation and centralization, what a top-level domain is and how it relates to our use of

[1] https://www.openintel.nl/

ccTLD. Whereafter we introduce the Simple Mail Transfer Protocol (SMTP), how email routing works and the `whois` protocol.

A. Consolidation and Centralization

According to the Internet Engineering Task Force (IETF), centralization is defined as "the ability of a single entity or small group to observe, capture, control or extract rent from the operation or use of an Internet function" [8]. Centralization can be seen as something undesirable, but sometimes beneficial. Opponents of centralization often name power imbalances, limited innovation, and monoculture among others as to why it is undesirable. On the other hand, centralization can be beneficial, for example in DNS, where a global source of truth is centralized such as the root zone of DNS [8].

In an Internet consolidation taxonomy, the IETF defined Internet consolidation as "the process of increasing control over Internet infrastructure and services by a small set of organisations" [9]. Centralization and consolidation are used interchangeably in the context of the Internet, yet they are different. While consolidation is primarily driven by economic factors, it can lead to a more centralized Internet architecture [10].

B. Top-level Domain

The Internet Corporation for Assigned Names and Numbers (ICANN) defines a Top-level domain (TLD) as "the right-most string, or series of characters, in every web address" [11]. In the case of `www.uva.nl`, this means that `.nl` is the right-most string and thus the top-level of this domain. TLDs are generally divided into two categories: generic ones and country-code ones, creating the abbreviations gTLD (Generic Top-level Domain) and ccTLD (Country-code Top-level Domain). Examples of gTLDs include: `.com`, `.info`, `.net`, and `.pro`, while ccTLDs include country-specific domains like `.nl` for the Netherlands, `.de` for Germany, and `.ru` for Russia. In email systems, ccTLDs can reflect localized email infrastructure, this means that we can deduct where the provider that hosts the email is located. It should be noted that not all TLDs are a single right-most string, other examples of TLDs include `.lib.ny.us` or `.otama.fukushima.jp`. To make sure we are not just isolating the word after the last dot, we make use of the public suffix list[2] as our guide to what TLDs we can encounter.

C. Simple Mail Transfer Protocol

The original SMTP was proposed in 1981, it outlined an ecosystem of servers that could be used to send, receive and relay mail [12]. The protocol works by establishing a connection and exchanging a few key parameters such as `EHLO`, `FROM`, `RCPT To`, followed by the actual email message itself.

[2]https://publicsuffix.org/

Later improvements to the protocol, such as SMTP Secure with STARTTLS implemented authentication and opportunistic encryption of messages. These days, more and more email servers use forced encryption or DNS-Based Authentication of Named Entities (DANE) to securely upgrade an unencrypted session to an encrypted one[3].

D. Mail Exchanger Record

The Mail Exchanger (MX) record is a type of DNS record that specifies a domain name and a preference value to which to send mail to. This domain name points to an Address (A) or AAAA record in a zone, which is the address of the receiving mail server. An SMTP mailer will query this record type at a DNS server to get a list of resource records, and attempt to send mail to the domain name with the lowest priority number, which is the highest preference. If this fails in a scenario with multiple resource records the mailer will attempt to send it to the domain name with the second-lowest priority, thus the second-highest preference [13]. This mechanism allows for redundancy and failover in email delivery.

E. WHOIS

The WHOIS protocol is a query/response protocol used to provide information services to internet users [14], this old protocol has been subject to online scrutiny over the last few years, with important internet services such as Certificate Authorities terminating `WHOIS` based certificate validation [15]–[17]. The main scrutinized elements are its old age and varying text-based formats lacking any form of uniformity, standardization or encryption. ICANN's proposed amendments will see the phase-out of this protocol starting in 2025.

III. Related Work

Internet centralization has been a discussion point for some time, with several studies done on this phenomenon. In the past, the DNS, hosting and email sectors have all been subject to such research. Zembruzki *et al.* [5] conducted a study concerning centralization within web hosting, and found that over 30% of all the domains are hosted by only 5, US-based, companies. They conducted this study by using large-scale DNS measurements provided to them by OpenINTEL and used the IP address pointed towards by an A or AAAA record to attribute a domain's hosting to an Autonomous System (AS) number.

In a paper authored by the Dutch domain registry, in collaboration with the New Zealand domain registry, Moura *et al.* [7], over 30% of DNS queries to both ccTLDs were also sent from 5 large cloud providers. Using historic DNS traffic data to each country's respective root servers, the researchers were able to map changes in the sources of queries over a period of 3 years.

[3]https://stats.dnssec-tools.org/#/?top=dane&trend_tab=1

A study conducted by Zembruzki *et al.* [6] found that the top 5 DNS providers account for more than 20% of all domains, and the top 100 providers for more than 80% across several different TLDs. Using OpenINTEL data, Zembruzki used the available Name Server (NS) record's IP addresses to attribute them to a DNS provider using the CAIDA prefix-to-AS mapping dataset[4], which is common among related studies in this field. Using a 5-year historical dataset, analyzing trends from 2017 up until 2021. One of the findings is the difference between TLDs in the concentration of providers, with `.ee` and `.ca` being the outliers.

In a study specifically looking at email providers, Liu *et al.* [18] attempted to characterize mail service providers using DNS data. Using a two-pronged approach, by looking at the available MX records combined with the domains in the STARTTLS certificates and EHLO messages. They were able to characterize these providers by looking at the email domains. When applying this methodology to their datasets, they found a handful of big email providers dominating the respective markets.

Zembruzki *et al.* [19] found high concentrations of recurring email providers within the analyzed TLDs, finding some TLDs with a top 5 dominance of up to 95%. This study uses a set of DNS data from the Tranco list. Each domain's MX record is resolved to determine their respective Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) records, after which this is then mapped to an AS number using the CAIDA prefix-to-AS mapping dataset[4].

Although the aforementioned work advances the knowledge about centralization in email, hosting and DNS, these studies lack a historical view of the European status quo in email hosting. The research on email lacks historical trends and does not consider Europe's current geopolitical situation. These works form a solid basis for our methodology, as well as give us an idea of what we can expect to see. We will contribute to the state-of-the-art by analyzing the trends and shifts in the email hosting landscape.

IV. Methodological Overview

In this section, we provide a detailed explanation of the methodology, its assumptions and limitations and the datasets used to study the centralization and consolidation of the email hosting industry.

A. Methodology

For our list of domains, we are using the zone-based DNS measurements from OpenINTEL because we want to create the most realistic and up-to-date measurements in terms of deployed domain names [20]. OpenINTEL extracts domain names from zone files of specific TLDs to create "largely complete" [21] measurements of the zones. In addition, OpenINTEL supports a wide range of DNS record types such as the required MX record type to indicate Mail Exchangers.

[4]https://www.caida.org/catalog/datasets/routeviews-prefix2as/

Earlier publications about email hosting and centralization make use of the aforementioned MX record as a key indicator for email hosting [18], [19]. Given the ease of using this data, and the earlier works supporting this methodology, we are using this record type to derive where the email for a domain is hosted.

We will use the associated MX record for a domain name to extract the second-level domain from the record. This second-level domain is extracted and used to profile at which provider a domain's email is hosted. An example scenario would look like this: `uva.nl. MX 100 uva-nl.mail.protection.outlook.com`. We can see that `uva.nl` has a single MX record with a priority of 100 pointing to `uva-nl.mail.protection.outlook.com`. We extract the second-level domain from the record which combines `outlook` and `.com` to form `outlook.com`. With the second-level domain extracted, we can count the number of domains with the specific second-level domain in their MX records and attempt to group them based on the domain name. In certain cases, domain names will be grouped for a provider based on Open-source Intelligence (OSINT). These will be domain names that we were not able to attribute via WHOIS or Registry Data Access Protocol (RDAP) but need to be grouped to keep the data as unbiased as possible.

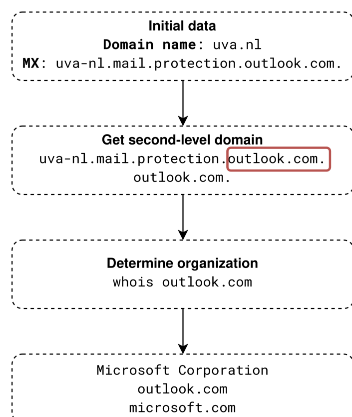

Fig. 1: Determine where mail is hosted

B. Assumptions and Limitations

In our research, we will encounter a few so-called "edge cases". In this section, we will describe them and explain how we deal with them. When registering a domain, one can choose to host content there or to reserve it for later use. The latter is also known as domain parking, in this case, a domain is registered at a domain registry but not actively used. When registering, a domain might be assigned the provider's default MX record. We do not make a distinction between actively used or parked domains if it has an MX record assigned it is included.

Year	State	`.se`	`.ee`	`.ch`	`.sk`	`.fr`
2025	Unique domains	1.39M	170K	2.53M	463k	4.16M
	Unique domains with at least 1 MX record	888K	121K	1.83M	361K	2.69M
	Percentage of domains with at least 1 MX record	63.6%	71.1%	72.3%	78.0%	64.8%
2024	Unique domains	1.41M	161K	2.52M	462k	4.16M
	Unique domains with at least 1 MX record	893K	118K	1.81M	354K	3.07M
	Percentage of domains with at least 1 MX record	63.0%	73.6%	72.0%	76.7%	73.9%
2023	Unique domains	1.43M	153K	2.47M	448k	4.16M
	Unique domains with at least 1 MX record	893K	111K	1.77M	351K	3.25M
	Percentage of domains with at least 1 MX record	62.2%	72.5%	71.6%	78.4%	78.3%
2022	Unique domains	1.40M	146K	2.34M	438k	4.03M
	Unique domains with at least 1 MX record	877K	107K	1.73M	342K	3.25M
	Percentage of domains with at least 1 MX record	62.6%	73.5%	74.0%	78.1%	80.8%
2021	Unique domains	1.51M	134K	2.21M	-	-
	Unique domains with at least 1 MX record	873K	100K	1.65M	-	-
	Percentage of domains with at least 1 MX record	57.8%	75.0%	74.9%	-	-
2020	Unique domains	1.46M	124K	2.13M	-	-
	Unique domains with at least 1 MX record	863K	95K	1.61M	-	-
	Percentage of domains with at least 1 MX record	59.1%	76.9%	75.3%	-	-

TABLE I: Number of unique domains and domains with at least one MX record per ccTLD

This means that we assume a domain has an email service if there is an associated MX record.

In some scenarios, a domain can have multiple MX records of varying priorities, in an attempt to attribute an email hoster to a domain's MX record, the lowest priority value will be used since the lowest value is the first domain name a mailer will try. In a scenario with multiple MX records with the same priority, a distinction is made between records with the same second-level domain and records with different second-level domains. When two identical second-level domains have the same priority, only one will be included. But when two unidentical second-level domains have the same priority, they will be included as two distinct entries.

C. Datasets

To measure the evolution of centralization in email hosting, we use historical data from OpenINTEL for the 1st day of each month covering the last 6 years (2020-2025). Table I gives an overview of the number of domains per ccTLD per year. For each ccTLD, we analyze the data of January 1st of the selected year. However, the dataset does not contain data for the `.fr` and `.sk` zones for 2020 and 2021, this is because the measurements for these zones only started in mid-2022 for both of these zones. This means that for both zones we take the earliest available date of 2022, which is 2022-08-10 for the `.fr` zone and 2022-05-11 for the `.sk` zone. This means that a complete comparison between all five zones is not possible, but the ability to spot certain trends persists.

D. Extrapolation

The striped bars in Figure 4 are based on estimated data since we do not have actual values for these years. To approximate these values, we use later years' data to derive past trends.

We first calculate the annual growth rate using:

$$\lambda_i = \frac{TLD_{i+1} - TLD_i}{TLD_i}$$

Then, we estimate past values by applying:

$$TLD_{i-1} = \frac{TLD_i}{1 + \lambda_{i-1}}$$

We calculate the value for 2021 using the growth rate from 2022, and then use this estimated 2021 value to calculate 2020.

V. Results

This section presents the outcomes of the study, which investigated the degree of consolidation and centralization of email hosting in European ccTLDs.

A. Trends

Our research findings indicate that there is an ongoing trend of centralization. Our aforementioned definition of centralization captures what we see throughout all five examined ccTLDs. Looking at Figure 2, we observe an increase in market share for the top 1-5 providers, year-over-year. Right now in the `.ee` and `.fr` zones we see the top 5 providers with a market share of 70.1% and 66.0%, the highest across all studied ccTLDs, with Sweden's `.se` sitting at 61.0%, Switzerland's `.ch` at 42.3% and Slovakia's `.sk` at a 29.1% for the top 1-5. In addition, the share of the rest group of the providers is steadily decreasing every year. Similarly, the top 21–100 hosting providers for each ccTLD are experiencing a decline, though the extent varies from year to year. Lastly, the top 11-20 and top 6-10 are in decline in zones such as `.fr`, `.ch` and `.ee`, but in the `.sk` zone we see a more or less steady top 11-20 and perhaps even a growing top 6-10.

When looking at the data in Table I we can see that some zones peaked in earlier years and that the total

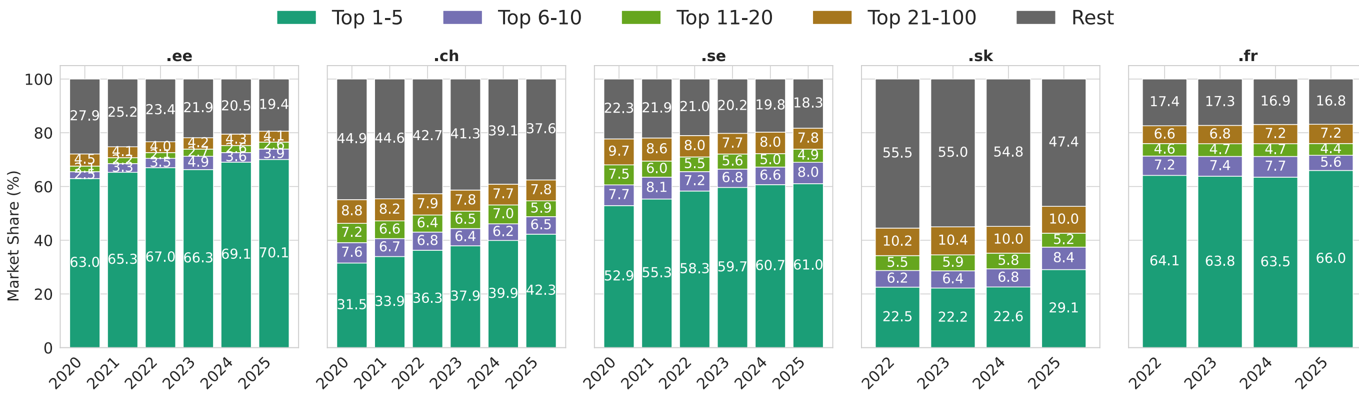

Fig. 2: Percentiles of providers per toplevel

number of unique domains is in slight decline, one such example is Sweden's `.se` zone. While the total number of domains has decreased, the degree of centralization keeps increasing. Growing the top 5's share from around 770K in 2020 to over 847K in 2025, while the total number of domains decreased by 70K.

Figure 4 shows us that across all examined ccTLDs, bigger pan-European and North American hosters are getting increasingly bigger. As companies and individuals leave smaller hosters, they move to bigger ones. We can observe Outlook sitting at just over 300K domains in 2022, while we continue to see growth in 2025 reaching 420K. If we exclude the data from the `sk` and `fr` zones, since they do not go back to 2020, we can see a sharp increase in Microsoft's Outlook share, almost doubling from 2020 to 2025.

B. Anomalies

In Table I we observe the raw number of unique domains in the data, and the number of unique domains with at least one MX record. During analysis, we found two anomalies, one of which we accounted for. The 01-01-2025 dataset for the `.sk` zone was missing 58K domains, as compared to 01-01-2024, this is a weird drop since the years before we saw an increase of about 10K total domains each year. Further analysis revealed that datasets for all dates after 01-01-2025 included a non-anomalous data point. For our comparison, we used the dataset from 02-01-2025, which contains 463K unique domains.

An additional inexplicable finding is that the Table I is the number of domains with at least one MX record decreased with over 360K between 2024 and 2025 for the `.fr` zone, re-runs on the dataset for other dates past 01-01-2025 yield the same result. From 2023 to 2024 there was also a decrease of 180K domains with at least one MX record, while the total number of domains is constant at 4.16M across 2023-2025.

C. Digital Sovereignty

When considering European digital sovereignty in light of our findings, we observe that Europe maintains a strong hosting sector. In each ccTLD studied, the top provider is consistently a local one from the respective country. This suggests that, at least at the domain level, European providers continue to play a dominant role. But we do have to consider the growing share of American companies. As shown in Figure 4, Microsoft's Outlook continues to grow annually across the zones, while Google's growth, though slower, remains steady. Other non-European email providers include Zoho, Godaddy and others.

While our study primarily focused on MX records of domains, it is important to acknowledge that some European providers may still rely on cloud services from providers such as AWS, Google Cloud, Microsoft Azure or others. This raises questions about the true extent of centralization and of the perceived digital sovereignty in this study. Further research could explore the degree to which European email providers depend on foreign cloud platforms, as this could reveal additional centralization concerns beyond domain-level hosting.

Additionally, we observed 8065 domains in the `.ch` zone that point to `example.org`, a non-existent MX record in Internet Assigned Numbers Authority's (IANA) system. This could indicate a dummy MX record or a standard configuration by a hoster that does not have or serve Email.

VI. Discussion

Among the top providers for a given ccTLD, we find that countries or regions with a shared language often have a dominant local provider or several smaller providers with similar market shares. Further down the list, historical relationships and ties between countries appear to be subtly reflected in our email hosting data.

For example: in Switzerland where in addition to German, French and Italian are spoken [22], we observe French hosting providers `gandi.net` and `OVH` being used, Italian provider `register.it` being used and German providers such as `ionos.de`, `ispgateway.de` or `udag.de`. A similar argument can be made for the `.fr` top-level domain, in which we can see strong ties to German hosters, as well as a very small fraction using Belgian ones. The `.se` top-level also shows connections to geographically proximate countries, with hosters such as `dandomain.dk`, `domeneshop.no`, `simply.com` and `one.com` showing ties to Norse and Danish hosting- and email providers. As well as the German `rzone.de` and `udag.de`.

The `.ee` and `.sk` zones both show interesting linguistic and historical ties, with Estonia's top 30 providers including the Russian `mail.ru` and `yandex.net`, which could be influenced by the significant Russian population in the country [23]. In addition, we see them making use of German and Lithuanian providers. Slovakia's `.sk` TLD is dominated by four major local hosting providers, including the Slovak-Czech Webglobe. Similarly, we see more Czech hosters recurring in the top 30 of this ccTLD, which could be because the Czech and Slovak languages are mutually intelligible or the fact that the two countries are neighbouring.

Our methodology stated that we would not differentiate between parked domains or not, this means that our data includes domains with MX records that are not meant to receive mail. One such example could be `bounce.domain.tld`, or the `example.org` case from earlier. Because we did not check if a certain domain can receive mail the numbers we present could look different if we did apply such a check.

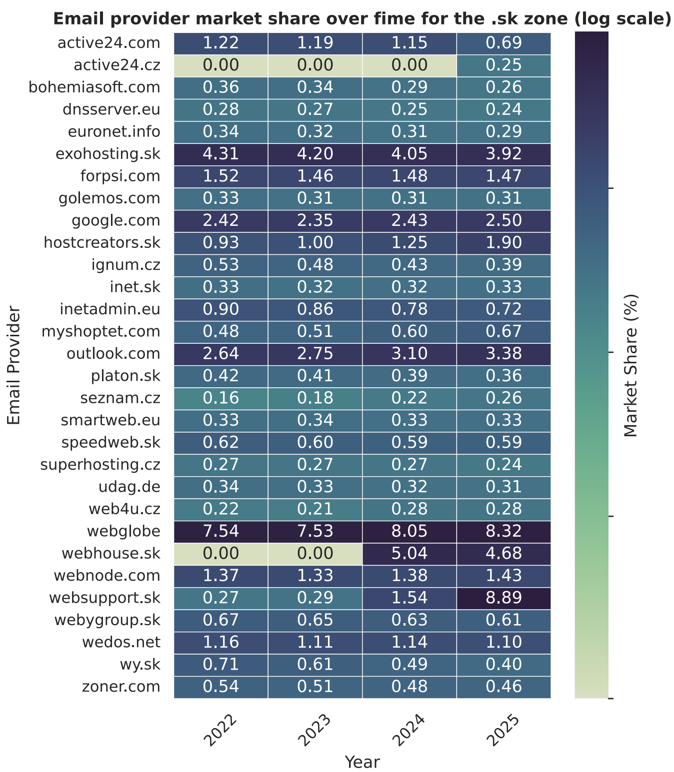

Fig. 3: This figure shows a heatmap of the providers in Slovakia's `.sk` zone and their percentage of total market share over the years on a logarithmic scale.

VII. Conclusions

In this paper, we investigated centralization trends within five European ccTLDs, finding that email hosting in these TLDs is increasingly concentrated among a few major providers. In some zones, such as Estonia's `.ee`, the top 5 providers collectively hold over 70% market share, while the smaller providers are slowly declining. However, centralization does not solely mean a shift from European to non-European providers. While Outlook has seen significant growth across all analyzed zones over the past six years, the overall landscape stays varied.

We found that the centralization of email hosting has evolved across the studied TLDs, and we observe a general trend of increased centralization. The degree of centralization and providers involved are less non-European than hypothesised, the implications for digital sovereignty are still significant seeing that domains are moving away from European hosters and migrating mostly to Outlook and Google Mail. These providers account for between 5 and 18% market share together, depending on the ccTLD.

Even though there is a strong European provider market which fills the rest, this does not automatically mean the mail is hosted in Europe. It could still be that these European providers use cloud services from providers outside of Europe to deliver their services, creating a distorted image of the state of centralization. The observed trends do highlight the challenge of keeping Europe's email ecosystem diverse and competitive.

VIII. Future Work

The ranking in this paper is purely based on domain count, the more a specific hoster occurs in the researched zone, the higher the perceived market share. What this method fails to take into consideration is the actual traffic to these MX servers and the actual hoster of these MX servers. Future research into email centralization could aim to measure the perceived top providers, in combination with how much traffic is being sent. With this setup, one can try to determine the biggest hosters in terms of traffic, and use the domain counts to enrich this data. This method would also partly filter our parked domains or domains that are generally not operational.

Another approach for future work could be a broader analysis of all OpenINTEL's ccTLDs, enriched by a historic DNS dataset. As this was our initial idea, but did not have the time or contacts to realize this, this can still be realized for a more in-depth analysis. This would also suit the current related works in centralization research since they follow a similar approach.

IX. Acknowledgements

I want to thank Moritz Müller from SIDN Labs for reviewing this paper and providing valuable feedback.

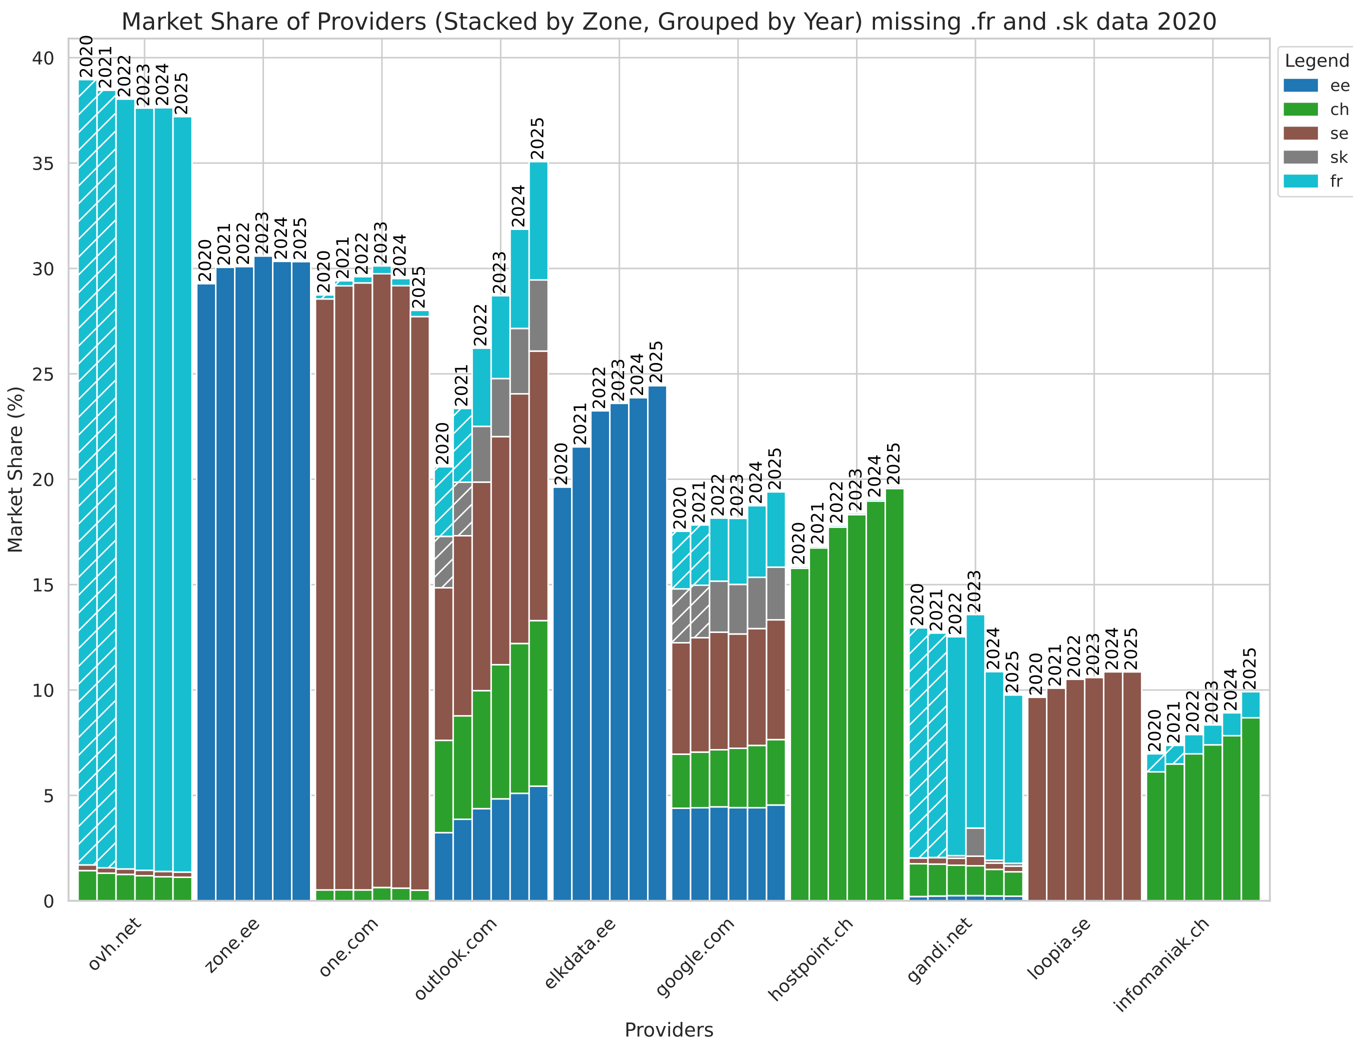

Fig. 4: For this figure, the top email hosters of every ccTLD were determined based on their size compared to the market share of the hoster, which are added to the bar chart as a stack to demonstrate relative size between hosters and between zones without clouding it with absolute values. The striped bars are based on extrapolated data as explained in IV.

References

[1] J. Arkko, B. Trammell, M. Nottingham, C. Huitema, M. Thomson, J. Tantsura, and N. t. Oever, "Considerations on Internet Consolidation and the Internet Architecture," Internet Engineering Task Force, Internet Draft draft-arkko-iab-internet-consolidation-02, Jul. 2019, num Pages: 16. [Online]. Available: https://datatracker.ietf.org/doc/draft-arkko-iab-internet-consolidation

[2] J. Schellevis, "Onrust over gedeeltelijke verhuizing .nl-domeinen naar het Amerikaanse Amazon," Feb. 2024. [Online]. Available: https://nos.nl/artikel/2507035-onrust-over-gedeeltelijke-verhuizing-nl-domeinen-naar-het-amerikaanse-amazon

[3] "Simple Mail Transfer Protocol," RFC 788, Nov. 1981. [Online]. Available: https://www.rfc-editor.org/info/rfc788

[4] "Digital sovereignty: Europe less vulnerable politically and economical." [Online]. Available: https://www.tno.nl/en/newsroom/insights/2022/06/strengthening-digital-sovereignty-makes/

[5] L. Zembruzki, R. Sommese, L. Z. Granville, A. Selle Jacobs, M. Jonker, and G. C. M. Moura, "Hosting Industry Centralization and Consolidation," in *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*, Apr. 2022, pp. 1–9, iSSN: 2374-9709. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9789881

[6] L. Zembruzki, A. S. Jacobs, and L. Z. Granville, "On the Consolidation of the Internet Domain Name System," in *GLOBECOM 2022 - 2022 IEEE Global Communications Conference*, Dec. 2022, pp. 2122–2127, iSSN: 2576-6813. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10001425

[7] G. C. M. Moura, S. Castro, W. Hardaker, M. Wullink, and C. Hesselman, "Clouding up the Internet: how centralized is DNS traffic becoming?" in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC '20. New York, NY, USA: Association for Computing Machinery, Oct. 2020, pp. 42–49. [Online]. Available: https://dl.acm.org/doi/10.1145/3419394.3423625

[8] M. Nottingham, "Centralization, Decentralization, and Internet Standards," RFC 9518, Dec. 2023. [Online]. Available: https://www.rfc-editor.org/info/rfc9518

[9] M. McFadden, "A Taxonomy of Internet Consolidation," Internet Engineering Task Force, Internet Draft draft-mcfadden-consolidation-taxonomy-03, Oct. 2023, num Pages:

12. [Online]. Available: https://datatracker.ietf.org/doc/draft-mcfadden-consolidation-taxonomy-03

[10] M. McFadden and D. Lazanski, "On the Effects of Internet Consolidation," Internet Engineering Task Force, Internet Draft draft-mcfadden-cnsldtn-effects-02, Oct. 2023, num Pages: 19. [Online]. Available: https://datatracker.ietf.org/doc/draft-mcfadden-cnsldtn-effects-02

[11] "Top-Level Domain." [Online]. Available: https://icannwiki.org/Top-Level_Domain

[12] "Simple Mail Transfer Protocol," RFC 821, Aug. 1982. [Online]. Available: https://www.rfc-editor.org/info/rfc821

[13] C. Partridge, "Mail routing and the domain system," Internet Engineering Task Force, Request for Comments RFC 974, Jan. 1986, num Pages: 7. [Online]. Available: https://datatracker.ietf.org/doc/rfc974

[14] L. Daigle, "WHOIS Protocol Specification," RFC 3912, Sep. 2004. [Online]. Available: https://www.rfc-editor.org/info/rfc3912

[15] "Old Protocol - whodis?" [Online]. Available: https://www.netmeister.org/blog/whois.html

[16] "We Spent $20 To Achieve RCE And Accidentally Became The Admins Of .MOBI," Sep. 2024. [Online]. Available: https://labs.watchtowr.com/we-spent-20-to-achieve-rce-and-accidentally-became-the-admins-of-mobi/

[17] "End of life for WHOIS-based DCV methods." [Online]. Available: https://knowledge.digicert.com/alerts/end-of-life-for-whois-based-email-dcv-method

[18] E. Liu, G. Akiwate, M. Jonker, A. Mirian, S. Savage, and G. M. Voelker, "Who's got your mail?: characterizing mail service provider usage," in *Proceedings of the 21st ACM Internet Measurement Conference*. Virtual Event: ACM, Nov. 2021, pp. 122–136. [Online]. Available: https://dl.acm.org/doi/10.1145/3487552.3487820

[19] L. Zembruzki, A. S. Jacobs, L. Z. Granville, and R. J. Pfitscher, "Examining the Centralization of Email Industry: A Landscape Analysis for IPv4 and IPv6," in *2023 IEEE Symposium on Computers and Communications (ISCC)*, Jul. 2023, pp. 360–366, iSSN: 2642-7389. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10217834

[20] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras, "A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 6, pp. 1877–1888, Jun. 2016, conference Name: IEEE Journal on Selected Areas in Communications. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7460220

[21] "OpenINTEL: Active DNS Measurement Project." [Online]. Available: https://openintel.nl/data/forward-dns/zone-based/

[22] "Main languages, 1970–2022." [Online]. Available: https://datawrapper.dwcdn.net/0702506434fb5c0560b5edbbe3d048a9/8/

[23] "RV0222U: POPULATION, 1 JANUARY by Year and Ethnic nationality. Statistical database." [Online]. Available: https://andmed.stat.ee/en/stat/rahvastik__rahvastikunaitajad-ja-koosseis___rahvaarv-ja-rahvastiku-koosseis/RV0222U/table/tableViewLayout2

Appendix A
Hosting providers top 10 for each ccTLD

ccTLD	Rank	2020	2021	2022	2023	2024	2025
.ee	1	zone.ee (28.088)	zone.ee (30.289)	zone.ee (32.327)	zone.ee (33.972)	zone.ee (35.945)	zone.ee (36.771)
	2	elkdata.ee (18.828)	elkdata.ee (21.702)	elkdata.ee (24.976)	elkdata.ee (26.211)	elkdata.ee (28.280)	elkdata.ee (29.634)
	3	elion.ee (4.906)	elion.ee (4.632)	google.com (4.788)	outlook.com (5.371)	outlook.com (6.040)	outlook.com (6.595)
	4	google.com (4.213)	google.com (4.448)	outlook.com (4.697)	google.com (4.909)	google.com (5.233)	google.com (5.511)
	5	outlook.com (3.101)	outlook.com (3.897)	elion.ee (4.284)	elion.ee (2.185)	telia.ee (3.449)	telia.ee (3.302)
	6	zoho (619)	serveriai.lt (912)	serveriai.lt (1.181)	telia.ee (2.034)	serveriai.lt (1.387)	serveriai.lt (1.420)
	7	yandex.net (547)	zoho (698)	zoho (760)	serveriai.lt (1.323)	zoho (948)	netim.net (1.114)
	8	netim.net (467)	yandex.net (665)	yandex.net (656)	zoho (846)	netim.net (763)	zoho (1.074)
	9	cscdns.net (409)	netim.net (534)	netim.net (627)	netim.net (669)	elion.ee (535)	hostinger.com (478)
	10	almic.ee (348)	cscdns.net (449)	cscdns.net (457)	yandex.net (543)	cscdns.net (454)	cscdns.net (445)
.ch	1	hostpoint.ch (254.224)	hostpoint.ch (277.431)	hostpoint.ch (308.310)	hostpoint.ch (324.708)	hostpoint.ch (344.449)	hostpoint.ch (357.747)
	2	infomaniak.ch (98.612)	infomaniak.ch (107.546)	infomaniak.ch (121.251)	infomaniak.ch (131.109)	infomaniak.ch (142.259)	infomaniak.ch (158.941)
	3	outlook.com (70.589)	outlook.com (81.328)	outlook.com (97.349)	outlook.com (112.801)	outlook.com (129.120)	outlook.com (143.781)
	4	jimdo.com (43.711)	google.com (43.767)	google.com (47.123)	google.com (49.908)	google.com (53.671)	google.com (56.810)
	5	google.com (41.340)	jimdo.com (41.968)	jimdo.com (38.319)	jimdo.com (34.956)	jimdo.com (31.736)	jimdo.com (28.454)
	6	pickelhost.com (30.301)	gandi.net (25.112)	gandi.net (25.089)	gandi.net (25.067)	servicehoster.ch (23.477)	servicehoster.ch (27.441)
	7	gandi.net (25.062)	hostcenter.com (23.424)	hostcenter.com (23.136)	servicehoster.ch (22.953)	gandi.net (22.927)	swizzonic.ch (25.528)
	8	hostcenter.com (23.267)	ovh.net (21.729)	servicehoster.ch (23.047)	ovh.net (21.071)	swizzonic.ch (21.547)	gandi.net (21.179)
	9	ovh.net (23.075)	tophost.ch (20.841)	ovh.net (21.779)	hostcenter.com (20.897)	ovh.net (20.686)	ovh.net (20.462)
	10	tophost.ch (20.828)	switchplus-mail.ch (17.716)	tophost.ch (21.074)	tophost.ch (20.797)	tophost.ch (20.225)	ionos.de (20.419)
.se	1	one.com (242.010)	one.com (250.390)	one.com (252.734)	one.com (260.145)	one.com (255.415)	one.com (241.715)
	2	loopia.se (83.322)	loopia.se (88.112)	loopia.se (92.215)	outlook.com (96.689)	outlook.com (105.837)	outlook.com (113.583)
	3	outlook.com (62.469)	outlook.com (74.651)	outlook.com (86.832)	loopia.se (94.620)	loopia.se (97.044)	loopia.se (96.442)
	4	google.com (45.615)	google.com (47.430)	google.com (48.961)	google.com (48.445)	google.com (49.527)	google.com (50.470)
	5	glesys.se (19.561)	glesys.se (18.797)	oderland.com (21.150)	oderland.com (22.815)	oderland.com (22.948)	oderland.com (23.776)
	6	fsdata.se (16.020)	oderland.com (18.603)	simply.com (18.762)	simply.com (19.689)	simply.com (20.519)	simply.com (20.141)
	7	oderland.com (16.005)	simply.com (15.637)	glesys.se (17.355)	glesys.se (16.600)	glesys.se (15.426)	rzone.de (15.266)
	8	misshosting.com (13.829)	h-email.net (14.504)	ilait.se (9.145)	ilait.se (8.953)	ilait.se (8.932)	glesys.se (14.007)
	9	h-email.net (10.112)	fsdata.se (11.254)	misshosting.com (8.757)	misshosting.com (7.747)	rzone.de (6.375)	websupport.se (9.944)
	10	surftown.se (10.099)	misshosting.com (10.541)	fsdata.se (8.143)	telia.com (6.903)	telia.com (6.159)	ilait.se (9.192)
.sk	1	—	—	webglobe (25.854)	webglobe (26.463)	webglobe (28.531)	websupport.sk (32.135)
	2	—	—	hostmaster.sk (19.595)	hostmaster.sk (19.191)	webhouse.sk (17.862)	webglobe (30.070)
	3	—	—	exohosting.sk (14.781)	exohosting.sk (14.780)	exohosting.sk (14.375)	webhouse.sk (16.920)
	4	—	—	outlook.com (9.068)	outlook.com (9.684)	outlook.com (10.998)	exohosting.sk (14.164)
	5	—	—	google.com (8.294)	google.com (8.272)	google.com (8.624)	outlook.com (12.213)
	6	—	—	forpsi.com (5.206)	forpsi.com (5.130)	websupport.sk (5.443)	google.com (9.038)
	7	—	—	webnode.com (4.715)	webnode.com (4.687)	forpsi.com (5.251)	hostcreators.sk (6.855)
	8	—	—	active24.com (4.199)	gandi.net (4.681)	webnode.com (4.876)	forpsi.com (5.300)
	9	—	—	wedos.net (3.987)	active24.com (4.193)	hostcreators.sk (4.421)	webnode.com (5.159)
	10	—	—	hostcreators.sk (3.193)	wedos.net (3.903)	active24.com (4.065)	wedos.net (3.992)
.fr	1	—	—	ovh.net (1.190.069)	ovh.net (1.178.033)	ovh.net (1.115.080)	ovh.net (966.653)
	2	—	—	gandi.net (338.374)	gandi.net (330.031)	gandi.net (275.367)	ionos.fr (330.217)
	3	—	—	ionos.fr (232.882)	ionos.fr (243.626)	ionos.fr (239.989)	gandi.net (215.315)
	4	—	—	1and1.fr (190.369)	1and1.fr (181.734)	1and1.fr (159.756)	outlook.com (151.313)
	5	—	—	outlook.com (120.824)	outlook.com (128.012)	outlook.com (144.811)	google.com (96.244)
	6	—	—	google.com (97.463)	google.com (101.955)	google.com (104.526)	m2bp (51.317)
	7	—	—	securemail.pro (47.399)	securemail.pro (46.371)	m2bp (39.040)	infomaniak.ch (33.335)
	8	—	—	m2bp (32.958)	m2bp (34.158)	securemail.pro (38.621)	securemail.pro (30.050)
	9	—	—	infomaniak.ch (29.744)	infomaniak.ch (30.756)	infomaniak.ch (33.330)	orange.fr (18.560)
	10	—	—	hostedemail.com (25.398)	hostedemail.com (24.791)	orange.fr (19.670)	hostinger.com (17.149)

TABLE II: Top 10 Hosting Providers for Each ccTLD (2020–2025)

Appendix B
Source code

All source code is openly available via https://gitlab.os3.nl/tseijsener/rp2-nlnetlabs-email-centralization