Analyzing the use of renewable energy in Dutch web hosting through DNS measurement data

Max Jeltes
University of Twente
P.O. Box 217, 7500AE Enschede
The Netherlands
m.jeltes@student.utwente.nl

The rise in the number of websites on the Internet has led to an increase in the number of data centers that consume electricity, indirectly emitting CO2. It is unclear how many websites in the Netherlands are currently hosted using renewable energy. Increasing attention to this topic could lead to more companies and individuals choosing a green hosting provider, thereby reducing the impact on the climate. During the study, we examined a variety of datasets of domains in order to determine which of these domains are hosted with renewable energy. The datasets examined are the most visited domains globally hosted in the Netherlands, the most visited domains by Internet users in the Netherlands hosted globally and Dutch government domains regardless of where they are hosted. Data from the Green Web Foundation was used to verify whether a domain is hosted using renewable energy. The study found that of the world's most visited domains hosted in the Netherlands, 29% of the approximately 80,000 domains examined are hosted using renewable energy. The study also found that 53.1% of the 1718 Dutch government domains analyzed, are hosted using renewable energy. Finally, the study estimated CO2 emissions for the 100 most visited domains by Internet users in the Netherlands. It was found that the majority of these domains are hosted using renewable energy and that hosting these 100 domains in October 2023 caused an estimated 5,287 metric tons of CO2 emissions. This is equivalent to the CO2 emissions from the electricity consumption of 1,029 households for one year.

Keywords
Renewable energy, Green Web Foundation, domain, web hosting, Internet, DNS

1 INTRODUCTION

Since the emergence of the Internet, web hosting has grown into a multi-billion-dollar industry. Web hosting providers, typically in the form of data centers, enable businesses and individuals worldwide to access Internet services, ranging from search engines to video streaming services. This technological advancement has created numerous opportunities worldwide, but it has also raised concerns. To meet the increasing demand, data centers worldwide operate large infrastructures that consume an estimated 1.3% of the world's electricity consumption, according to [1]. Along with the data transmission networks required to run the Internet, this is estimated to have caused 330 Mt of CO2 emissions in 2020 [2]. According to the Dutch Central Bureau of Statistics [3], data centers accounted for 2.7% of total energy consumption in the Netherlands in 2019.

Although fossil fuels are predicted to remain the primary source of energy worldwide until 2035 [4], an increasing number of data centers are making efforts to become more sustainable. For instance, major data center operators like Google and Facebook are taking action to enhance fuel efficiency and reduce wasted power [5]. Various studies have been conducted on this topic, such as optimizing load balancing algorithms to intelligently cycle servers on and off based on demand [6], [7]. These initiatives aim to reduce energy consumption and, consequently CO2 emissions. However, data centers can also reduce their CO2 emissions by switching to green energy, which is generated from renewable and environmentally friendly sources. Green energy sources are sustainable and have a lower impact on the environment compared to traditional sources of energy, such as fossil fuels (coal, oil, and natural gas) [8]. Data centers can also become more sustainable through carbon offsetting, which involves compensating for produced CO2 emissions by financially investing in projects that reduce CO2 gases from the atmosphere [9].

This study utilized measurement data from the domain name system to obtain information on the number of hosting providers in the Netherlands that use green energy. The domain name system (DNS) translates domain names into IP addresses. This IP address can be used by a client on the Internet to connect to the corresponding server and request the webpage associated with this domain name. Historical DNS data comprises numerous records of domain names linked to their respective server IP addresses. IP-based geolocation can be used to determine the approximate location of the server, providing insight into the hosting location of a domain. DNS data can also potentially provide insight into the number of domains hosted by a hosting provider and the number of visitors those domains receive on a regular basis. However, using public DNS data for research presents challenges due to the high volume of daily Internet queries. Cloudflare, one of the largest cloud providers globally, processes over 2 trillion DNS queries daily [10]. Cloudflare's public DNS resolver, identifiable by the IPv4 address 1.1.1.1., is mainly responsible for handling these queries. Collecting DNS data is challenging due to the complex and large structure of the Internet. The processing and analysis of DNS measurements also require significant computing power. The OpenINTEL DNS

TScIT 40, February 2, 2024, Enschede, The Netherlands

measurement project [11], [24] is useful in this regard, capturing daily snapshots of the state of significant portions of the global DNS. The project focuses on a limited number of generic Top-Level Domains (TLD) and country code TLDs to determine which parts of the DNS are queried daily. OpenINTEL also queries the DNS based on lists of the most visited domains. These lists are based on the number of DNS queries observed at major DNS resolvers from companies such as Cloudflare and Amazon. Additional data, such as the location of a server hosting a particular domain, is retrieved via external APIs. This location may not be entirely accurate, but it provides enough detail to determine the country in which the server is located [12]. Additionally, the IP address of the queried domain name can be used to identify the autonomous system (AS) to which the server belongs. An AS is a large network or group of networks operated by a single organization, such as an Internet service provider (ISP) or a large enterprise technology company. Within an autonomous system, a single routing policy applies. This routing policy ensures that all devices within an AS can communicate with each other and reach devices in other ASs via the Internet. The Internet comprises a vast network of Autonomous Systems (ASs). The Border Gateway Protocol (BGP) was developed to route messages across the Internet and through the appropriate ASs. The BGP examines all available paths that data could take and selects the optimal route, typically involving multiple autonomous systems.

To determine whether a domain is hosted using green energy, a reference source is needed. The Green Web Foundation (GWF) maintains a public dataset [13]containing over 5 million domains for which the administrator/hosting provider has provided evidence that these domains are hosted using green energy. This proof must be provided on an annual basis to maintain the accuracy of the dataset. When a domain appears in GWF's dataset, it indicates that the hosting provider either uses 100% renewable energy or offsets the amount of CO2 emitted from non-renewable sources through carbon offsetting.

This paper is organized as follows. Section 2 provides a brief explanation of the problem and outlines the research questions. Section 3 summarises research related to the energy consumption of data centers and network infrastructure and its optimization. For each of the three sub-questions, Section 4 explains which dataset was used, how the data were analyzed and how the results were collected. Section 5 presents the study's results, which are divided by sub-question. Sections 6 and 7 present the study's results and limitations, respectively. Section 8 provides recommendations for future research.

2 PROBLEM STATEMENT

The demand for web hosting is increasing, leading to a growth in the number of data centers. This leads to an increase in energy consumption and CO2 emissions. Although some web hosting providers already operate using renewable energy, it is unclear what percentage of domains hosted in the Netherlands fall under this category. Conducting research based on DNS data can shed light on the current situation and the potential for reducing CO2 emissions.

2.1 Research question

Based on the problem statement, the following research question was formulated:

RQ To what extent do web hosting providers located in the Netherlands use renewable energy to make their websites available on the Internet according to the Green Web Foundation?

This question is answered using the following sub-questions:

SQ1 What percentage of the top 1 million most visited domains hosted in the Netherlands is powered by renewable energy?

SQ2 What percentage of domains hosted by the Dutch government run on renewable energy?

SQ3 What is the estimated carbon emission footprint of hosting the top 100 most visited domains in the Netherlands? Can the transition of non-certified hosting providers inside the top 100 to certified green energy hosting achieve any carbon emission reduction?

3 RELATED WORK

Several studies have been conducted on the energy consumption of data centers. For instance, Lü et al. [14] analyzed the energy consumption of a data center in Finland and found that air circulation was excessive. They suggested that a heat recovery system could be used to reuse the heat produced to warm other buildings. In another study [15], over 200 models were analyzed to estimate the energy consumption of a data center. The study revealed that the cooling system accounts for approximately 50% of the energy consumption in a data center, while servers and storage devices account for 26%. Additionally, the research found that in an average data center, IT equipment consumes approximately 30% of the energy, while cooling, lighting, network equipment, and emergency power supplies consume the remaining 70%.

Google's environmental report for 2023 [16] revealed that the company's total energy contracts in 2022 amounted to 11,600 MW worldwide, resulting in an energy consumption of 21,776,200 Gwh. To put this into perspective, this is equivalent to the energy consumption of approximately 2.05 billion US households, with an average annual energy consumption of 10.632 kWh. It is worth noting that the Bitcoin network is also a significant energy consumer. This study [17] estimated in 2017 that the global Bitcoin network requires between 100 and 500 MW of power.

In addition to mapping data center energy consumption, research was also carried out to optimize it. According to [18], building a data center near a water dam can result in energy savings of 12-15%. The water can be used for both electricity generation and cooling. Resource scheduling algorithms can also be used to switch off part of the server capacity during low demand, further optimizing energy consumption. Switching to (partly) renewable energy sources can reduce the operational cost and carbon footprint of a data center. However, switching data centers to renewable energy sources is proving challenging due to the variable yield of these sources and the uncertainty of their availability [19].

Finally, research has also been conducted on the energy consumption of network infrastructure and ways to reduce it. For example, in this [20] study, a clustering approach based on a spectral algorithm was used to reduce energy consumption by turning off part of the network when traffic is low. This algorithm has the potential to save significant amounts of energy

when implemented in large networks and large enterprises.

4 METHODOLOGIES

Python was used to analyze multiple datasets. A different set of domain names was used for each sub-question. For sub-question 1, we looked at the most visited domains worldwide hosted in the Netherlands based on the Cloudflare Radar top 1 million [21]. For sub-question 2, we looked at the Dutch government domains registered in the central government website registry [22]. And for sub-question 3, we looked at the 100 most visited domains by Internet users in the Netherlands based on data from Semrush [23], regardless of which country these domains are hosted from. To determine if a domain is hosted using renewable energy, the GWF dataset was used. In consultation with the organization, a list of AS numbers (ASN) from 22 September 2023 was provided with evidence that these ASs are operated using green energy. During the study, BGP data has been used to find out for each domain in which AS the hosting server is located. The Python package pyasn, created by Github user hadiasghari, was used for this lookup. The assumption was made that all servers located in a green hosted AS are themselves running on green energy.

To map a domain name to the data center where it is hosted, we rely on DNS and AS information, as well as the dataset provided by the GWF. The DNS is among others responsible for mapping domain names to IP addresses. When a client requests a web page, such as 'example.com', a DNS query is sent to a recursive resolver. The resolver then requests the IP address of an associated Top-Level Domain (TLD) server (in this case .com) from a root name server. The resolver sends a request to the TLD server, which responds with the IP address of the example.com nameserver. Finally, the recursive resolver sends a request to the domain's authoritative name server, which responds with the IP address of the requested domain name [25].

4.1 On answering SQ1

A list of domains and their corresponding ASs was created using an OpenINTEL snapshot [24] of the Cloudflare Radar top 1 million dataset. The snapshot includes DNS queries and responses for all domains in the Cloudflare top 1 million on September 22, 2023. Only the responses for DNS A records and DNS AAAA records were used. Each response includes a country code indicating the geographical location of the queried server. By filtering for the country code NL, all DNS responses have been obtained exclusively from servers located in the Netherlands. The reason for filtering on servers in the Netherlands is that the scope of the study was narrowed down to the Netherlands only. This limits the amount of DNS data to be analyzed and makes the analysis feasible. Each response includes an AS number indicating the location of the server on the Internet. The GWF's list of certified ASs is used to determine if a domain is hosted using green energy.

4.2 On answering SQ2

In order to answer this research question, the central government website register [22] was used. This register is a monthly publication that contains a list of domains managed by one of the 12 ministries within the Dutch government or by another government entity, such as the police or tax authorities. To exclude duplicates, only 2nd level domains (example.com) were included in the dataset, as the list also contains 3rd level domains (such as test.example.com) in which the 2nd level domain appears multiple times. To determine the location of the hosting server, the BGP was used as the list does not provide information about the AS of the hosting server. Once the AS number has been found, it was compared to the GWF dataset to determine if the server is hosted using green energy. The process of identifying green energy hosting based on the AS number is explained in detail at the beginning of this section.

4.3 On answering SQ3

Data from the SEO company Semrush was used to answer this question. The top 100 most visited domains by Internet users in the Netherlands [23] were compiled by Semrush by analyzing clickstream data. This data provides an aggregated view of the online journeys of millions of real but anonymized Internet users by tracking their online activity. The dataset provides a monthly estimate of the Internet behavior of people in the Netherlands, with the disclaimer that the estimate may differ from the actual situation. Besides the ranking of the top 100 domains, an estimated number of visitors per month is also available for each domain.

To estimate the total CO2 emissions caused by hosting the 100 most visited domains in the Netherlands, an average emission per domain is needed. For this purpose, the Website Carbon API [26], made by Wholegrain Digital was used. This API reports the estimated CO2 emissions for each domain accessible online, each time the corresponding website is viewed. This estimate is based on the amount of data sent over the Internet to a client's device to load the page. Using an estimated conversion factor of 0.81 kWh/GB, the page size can be converted to a given amount of energy consumed (kWh). This factor takes into account caching of the web page on both the client and server side. Finally, Wholegrain Digital is able to convert this amount of energy into a carbon footprint using two different estimation factors. One for a datacenter using standard grid energy (442g CO2 emissions/kWh) and the other for a datacenter using renewable energy (50g CO2 emissions/kWh). A detailed explanation of the calculation method used can be found in [27].

In addition to carbon emissions, several other metrics were collected for each domain to assess site quality. The amount of data required to load each page was recorded (in KB). The time it took to load the page in a browser was also collected (in ms). Finally, for each domain, the number of additional HTTP requests the page makes to retrieve additional information such as images and videos located on the web page. The above metrics were collected using the Pingdom tool [28] from SolarWinds. To provide insight into the calculated CO2 emissions, the U.S. Environmental Protection Agency's Greenhouse Gas Equivalencies Calculator [29] was used.

5 RESULTS

To determine if an autonomous system is certified according to the GWF, a dataset of verified ASs from the GWF was used. The dataset, created on 22 September 2023, contains 391 ASN's registered in 35 different countries. The majority of ASs (93) in the dataset are registered in the United States, followed by the Netherlands with 59 ASs and then Germany with 57 ASs. The dataset also indicates the company under which the AS is registered. The dataset includes 188 companies, with Equinix operating the highest number of ASs at 128, followed by Akamai Technologies with 26. The majority of companies have 1 or 2 ASs in the dataset.

5.1 Top 1 million domains

Geographical distribution

Since the Cloudflare Radar Top 1 million is published daily and querying the DNS for 1 million domains returns large amounts of data, it was decided to focus on the OpenINTEL DNS dataset for one day, September 22, 2023. Also available from the same day is a list of verified AS numbers from the GWF. Based on the analysis of DNS data, figure 1 shows per country how many domains from the Cloudflare Radar top 1 million are hosted in that country. The red countries on the map indicate that many domains are hosted in those countries. The blue countries have few domains hosted there. The map shows that many domains are hosted by servers located in the United States, Germany, and the Netherlands. Few domains are hosted in Africa. Regarding the ASs hosting the 1 million domains, Cloudflare's AS is by far the largest with 29.7% of the domains. Amazon's AS follows with 13.9%. Looking at the occurrence of different TLDs among the 1 million domains, .com is the most common with 42.5%. This is followed by .net with 8.5% and then .ru with 4.8%.

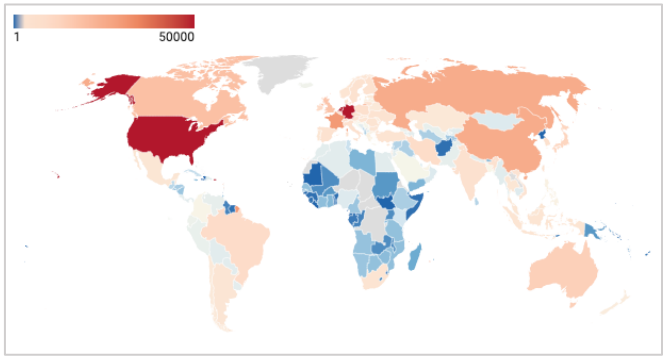

Figure 1: Number of hosted domains from Cloudflare's top 1 million radar per country.

Distribution of autonomous systems

After filtering the domains hosted on servers located in the Netherlands, 80,158 domains remained. Of these, 23,257 domains (29%) are hosted within an AS running on green energy. Among these AS, Akamai's AS is the largest with 10,982 domains. This is followed by Leaseweb with 2,348 domains and then TransIP with 1,499 domains. The other 56,901 domains (71%) are hosted in the Netherlands within an AS that is not certified by the GWF. This means that it is not known to the GWF what kind of energy the AS is running on. We therefore consider these domains as non-certified domains. Figure 2 shows the distribution of domains across the most popular non-certified ASs in the dataset. About a quarter of the non-certified domains in the dataset are hosted in Amazon's AS 16509. This is followed by Scalaxy's AS 58061 with 8.8% and then Digital Ocean's AS 14061 with 8.4%. The remaining 31.5% consists of ASs from which less than 2% of the non-certified domains in the dataset are hosted.

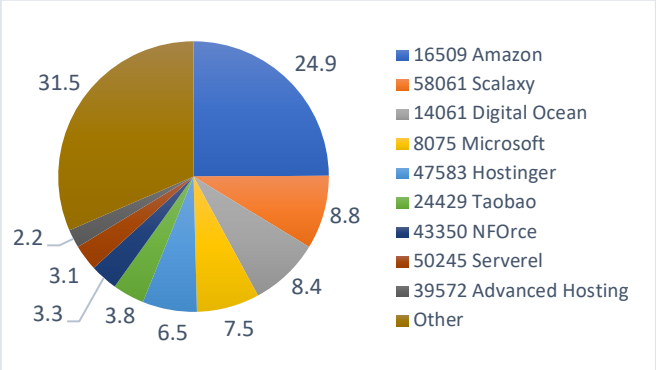

Figure 2: Distribution of Cloudflare Radar top 1 million domains across non-certified autonomous systems.

According to the GWF website, parts of Amazon's infrastructure are certified. The website also states that certification is dependent on the Amazon Web Services (AWS) region where the hosting server is located. This region is a physical location around the world where multiple Amazon data centers are clustered. Our investigation revealed that a quarter of the uncertified domains in the dataset are hosted from Amazon's AS 16509. We investigated which AWS regions these domains are hosted from. This showed that of the 14,164 domains in the dataset hosted from AS 16509, 90% were situated in the global region. This region is not tied to a specific location and is spread all over the world. The remaining 10% of the domains were hosted from the AWS regions eu-north-1, ap-south-1, and sa-east-1. These four regions are not included in the AWS regions certified by the GWF. This confirms that the domains found from Amazon's AS 16509 are not certified.

Top-level domains

Within the filtered dataset of domains hosted in the Netherlands, the set was analyzed to see which TLDs occurred. A total of 268 different TLDs were found. For each TLD, it was determined how many domains in the dataset had that TLD, and which of those domains were certified with the GWF and which were not. For 142 TLDs, it was found that the majority of domains with that TLD are hosted with green energy. For the other 126 TLDs, more domains per TLD were uncertified than certified. The most common TLDs are shown in figure 3. The figure also shows the number of certified and uncertified hosted domains for each TLD. There is a clear difference between TLDs where the majority of domains are certified (using green energy) (e.g. .nl and .de) and TLDs where the majority of domains are non-certified (e.g. .win and .xyz).

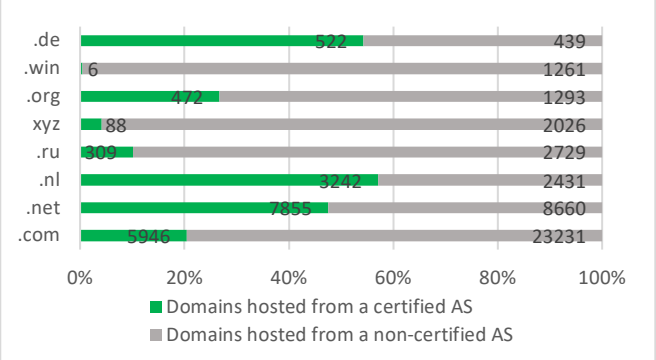

Figure 3: Distribution of certified and non-certified domains among the most common top-level domains.

5.2 Dutch government

For this sub-question, the dataset of the central government website registry was used. This dataset is updated and published on a monthly basis. The dataset obtained from the GWF was created on September 22, 2023. Unfortunately, since there is no September 2023 version of the website registry available, the August 2023, version was selected as the data input instead. Ideally, both datasets are from the same day, but a one-month difference is acceptable considering the website registry only adds a few domains per month. To avoid incorrect domain certification, it is crucial that the GWF dataset used is from the same day or a more recent date than the domains being checked.

Within the dataset, 1814 domain names were present. After filtering out the duplicate second level domains, 1718 domains remained. These domains fall into 22 different categories. Of these, 12 categories are ministries and the other 10 are government organizations such as the House of Representatives and the National Police. All of these organizations within the government own one or more domain name that they use to host websites for publishing information to citizens, businesses, and other entities. The majority (93.1%) of the domains contained the .nl TLD. Other TLDs that also appeared frequently are .com, .eu and .info. Grouping the domains by AS of the hosting server (see figure 4) shows that Prolocation has by far the most government domains in operation (26.7%). This is followed by some well-known Dutch ICT companies such as SURFnet (5.5%) and TransIP (3.9%). But also, global players like Microsoft and Amazon are present in the list and manage one or more domains for the Dutch government. The remaining 35.4% is made up of ASs from which less than 2% of the domains in the dataset are hosted.

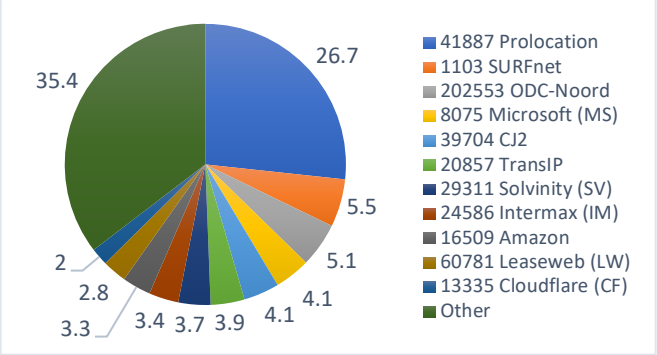

Figure 4: Distribution of Dutch government domains across autonomous systems.

The 1718 domains listed in the dataset are hosted by servers located in 113 different ASs. After verifying the AS numbers against the GWF dataset, it appears that the domains in the list are hosted by 32 different ASs that are certified, and therefore proven to use green energy or offset non-green energy. The remaining 81 ASs do not appear in the GWF dataset, and therefore the operators of these ASs have not provided evidence to the GWF that they use green energy. Figure 5 provides an overview of the largest certified and non-certified ASs and the number of domains that appeared on them. This shows that Prolocation has proven to the GWF that they are running their AS with green energy. However, there are still relatively large ASs managed by SURFnet and ODC-Noord that have not proven this. This could potentially mean they are not using green energy. The other option is that they are using green energy but have not reported this to the GWF.

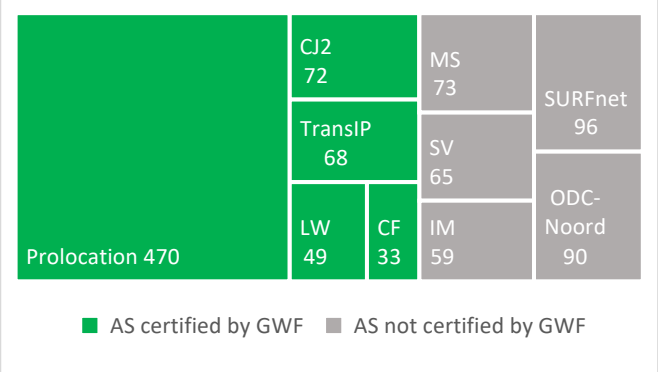

Figure 5: Overview of the largest certified and non-certified autonomous systems by number of hosted domains.

In total, 912 (53.1%) of the 1718 domains are hosted by servers in a certified green AS. The other 806 domains (46.9%) are located in non-certified AS. The vast majority (96.6%) of the domains are related to one of the 12 Dutch ministries. Figure 6 shows per ministry how many domains occur in the dataset and what percentage of them are hosted with green energy. The Ministry of General Affairs (AZ) stands out because almost all of its domains are hosted using green energy. The Ministry of Finance (FIN) also stands out as 66.1% of its domains are hosted from non-certified ASs. Out of 1718 domains, the dataset also includes the total number of visitors in August 2023 for 811 domains. The website www.rijksoverheid.nl received the most traffic with 10,956,312 visitors. The average number of visitors per domain in August 2023 was 35,400 visitors. Of the 811 domains for which the number of visitors in August 2023 was known, 69% were found to be hosted using green energy. Within both the top 10 and top 100 most visited domains, 60% of domains were hosted using green energy. The average number of visitors of all green-hosted domains was 40,000 visitors. Of the non-certified domains, the average number of visitors was 28,000 visitors.

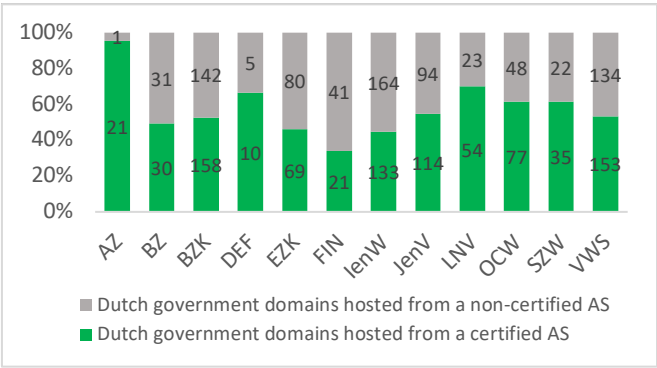

Figure 6: Overview of the number of certified and non-certified domains per Dutch government ministry.

5.3 Carbon emission footprint

The list of the top 100 most visited domains by Internet users in the Netherlands was used as a data source. This list of domains was published by Semrush and also includes for each domain the estimated number of visitors the domain received in October 2023. Additional information was then collected for each

domain, such as page size, load time, number of additional requests, and whether the domain was hosted from a certified AS. The number of visitors per domain ranged from 11.5 million (indeed.com) to 2.07 billion (google.com). The average number of visitors in October 2023 across the 100 domains was 778 million. The number of additional HTTP requests required to load a domain's web page ranged from 0 to 351 requests, with an average of 34 requests. Furthermore, the time taken to load the entire web page associated with the domain ranged from 31 to 9920ms, averaging 526ms. Moreover, the size of the loaded web page ranged from 80KB to 6.6MB with an average of 2.6MB. Finally, the estimated CO2 emissions per domain visit ranged from 0.15g to 2.09g with an average of 0.82g. These CO2 emissions were calculated by Wholegrain Digital [26] based on the amount of data sent over the Internet to a client's device to load the page. Figure 7 plots, for all 100 domains, the size of each web page loaded against the estimated CO2 emissions released in the process. This clearly shows that a larger page size results in more CO2 emissions per page load.

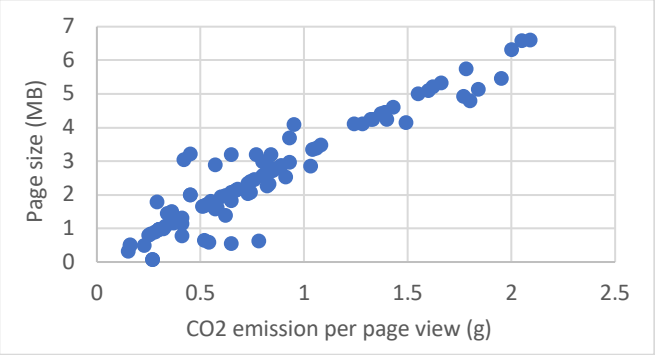

Figure 7: Page size compared to CO2 emissions per page view of top 100 most visited domains in the Netherlands.

Estimated carbon emission

Given the estimated CO2 emissions per page visit for each of the 100 domains and an estimate of the number of visitors per domain in October 2023, the total amount of CO2 emitted can be calculated. Multiplying the number of visitors for each of the 100 domains by the corresponding estimated CO2 emissions per page visit and adding these values together yields 5.287 million kg of CO2. This means that hosting the top 100 most visited domains in the Netherlands in October is estimated to have caused 5,287 metric tons of CO2 emissions. This is equivalent to the CO2 emissions from 2.25 million liters of gasoline consumed, or the electricity consumption of 1,029 households for one year. To offset the same amount of CO2 emissions, you would have to plant 87,421 trees and let them grow for 10 years.

Possible carbon emission reduction

Of the 100 domains, 83 were found to be hosted on a server located in a certified AS. Collectively, these 83 green domains caused an estimated 4543 metric tons of CO2 emissions in October 2023. The remaining 17 domains generated an estimated 744 metric tons of CO2 emissions. According to Wholegrain Digital, switching to green hosting could save 9% of CO2 emissions. Assuming that the 17 domains are currently hosted using standard grid energy, switching to green energy could save 67 metric tons of CO2 emissions. This CO2 reduction is 1.3% of the total estimated CO2 emissions released from hosting all 100 domains. This is equivalent to the CO2 emissions from 28,538 liters of gasoline consumed, or the electricity used by 13 households for one year.

5.4 Summary

The study found that of the world's 1 million most visited domains, just over 80,000 are hosted in the Netherlands. Of these, 29% were found to be hosted using green energy. The largest non-certified AS was Amazon's, with nearly a quarter of the domains. It also revealed that of the 1718 Dutch government domains examined, 53.1% are hosted using green energy. Finally, the top 100 most visited domains in the Netherlands were found to have caused an estimated 5,287 metric tons of CO2 emissions in October 2023. This is equivalent to the CO2 emissions from 2.25 million liters of gasoline consumed, or the electricity consumption of 1,029 households for one year.

6 DISCUSSION

Considering SQ 1, our expectation was that a large proportion of the 80,000 examined domains which are hosted in the Netherlands, would be hosted using renewable energy. This is due to research [30] indicating that, in August 2020, of the 475 green hosting providers present in the GWF dataset, 177 were located in the Netherlands. Germany came in second place with only 40 green hosting providers, despite having a population approximately five times larger than the Netherlands. At that moment, it appeared that the Netherlands was a frontrunner in green web hosting. Furthermore, we hoped that a company as large as Amazon would set a good example, however it was found that their AS is not certified according to the GWF. Upon further investigation, it was discovered that almost all domains within Amazon's AS are hosted in the global AWS region, which is according to the GWF not certified. It is possible that domains with high global traffic are hosted in this region to reduce worldwide latency.

For SQ 2, we expected the majority of Dutch government domains to be hosted using renewable energy, as the Dutch government is a proponent of using renewable energy and setting a good example is crucial. With 53.1% green-hosted domains, this expectation was just met. Worth noting is that the majority of domains under the Ministry of Economic Affairs and Climate are not certified, despite the fact that the ministry promotes clean energy [31].

The results of SQ 3 differ from those of SQ1. In SQ 1, 29% of the domains hosted in the Netherlands were certified, while in SQ 3, 83% of the examined domains were certified. One possible explanation for this difference is that the dataset of SQ 1 is much larger (80,000 domains) than in SQ 3 (100 domains). Therefore, it is still possible that within the 100 most visited domains out of the 80,000 domains examined from SQ 1, there were a similar number of certified domains as in SQ 3. However, it was not possible to test this as the SQ 1 dataset did not include a ranking based on the number of visitors per domain, whereas the SQ 3 dataset did.

6.1 Limitations

The study has a number of limitations. Firstly, for ASs that appear in the datasets studied but not in the GWF dataset, the only conclusion that can be drawn is that they were not certified at the time of measurement. However, this does not indicate the actual energy source of the AS, which could be either green or

not green. Therefore, the actual number of green-hosted domains may be higher than reported. Secondly, for sub-question 3, a list of the 100 most visited domains in the Netherlands was used. This list was compiled by Semrush using clickstream data collected by the company. It should be noted that the ranking and visitor counts in the list may differ from the actual numbers. To obtain a more accurate list, DNS data from major DNS resolvers such as Cloudflare could be used. Although these types of datasets are available online, they do not provide an estimated number of monthly visitors per domain. Finally, for sub-question 3, an estimate was made of the CO2 emitted by hosting the 100 most visited domains by Internet users in the Netherlands in October 2023. To make this estimation, Wholegrain Digital's Website Carbon API was used. The estimation includes four types of CO2 emissions: client device usage, network use, data center use, and hardware production. Wholegrain also uses a separate carbon intensity factor for each country. Despite this comprehensive calculation method, the actual amount of CO2 emitted can differ from the calculated amount. This is because it is difficult to make a precise estimate of the CO2 released when a web page is loaded. To do this, you would need to take into account the 4 types of CO2 emissions mentioned earlier. And each of these factors can vary widely, such as the network infrastructure, which in turn affects the CO2 emissions of that part of the chain.

7 CONCLUSION

The study's main question is: “To what extent do web hosting providers located in the Netherlands use renewable energy to make their websites available on the Internet according to the Green Web Foundation?”.

Regarding SQ 1, it can be concluded that a minority (29%) of the top 1 million most visited domains worldwide, hosted in the Netherlands, are certified according to the GWF. According to the results of SQ 2, we can conclude that 53.1% of the Dutch government domains are hosted using renewable energy. Switching AS operators, such as SURFnet, that manage many Dutch government domains to green energy would help to increase this number. To answer SQ 3, the estimated CO2 emissions for hosting the top 100 most visited domains in the Netherlands is 5,287 metric tons. This is equivalent to the CO2 emissions from the electricity consumption of 1,029 households for one year. If we assume that the 17 non-certified domains in this top 100 are currently hosted with standard grid energy, then switching to green hosting for these domains would have resulted in CO2 savings of 1.3% of total emissions in October 2023.

To answer the main question, we examined several datasets, each with a different perspective. In terms of the most visited domains worldwide, hosted in the Netherlands, 29% of green-hosted domains may actually be quite significant compared to other EU countries [30]. Upon examining the Dutch government, it can be concluded that slightly more than 50% of registered domains in August 2023 are hosted using green energy. Compared to the result of sub-question 1, this percentage can be considered significant.

8 RECOMMENDATIONS

Further research using data from previous years could provide a picture of the increase or decrease in the number of green-hosted domains in the Netherlands. Additional research could also be conducted on other sectors in the Netherlands, such as finance, healthcare, and education, to determine the extent to which related domains are hosted using renewable energy. Finally, it would be valuable to research the regulation of green web hosting by the Dutch government to determine whether legislation exists or could be proposed. This might push large hosting companies to invest more in sustainable hosting.

REFERENCES

[1] R. Bose, S. Roy, H. Mondal, D. R. Chowdhury, and S. Chakraborty, "Energy-efficient approach to lower the carbon emissions of data centers," *Computing*, vol. 103, no. 8, pp. 1703–1721, Aug. 2021, doi: 10.1007/S00607-020-00889-4/FIGURES/13.

[2] "Data centres & networks - IEA." Accessed: Nov. 20, 2023. [Online]. Available: https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks

[3] "Elektriciteit geleverd aan datacenters, 2017-2019." Accessed: Nov. 24, 2023. [Online]. Available: https://www.cbs.nl/nl-nl/maatwerk/2020/51/elektriciteit-geleverd-aan-datacenters-2017-2019

[4] A. Jäger-Waldau, M. Szabó, N. Scarlat, and F. Monforti-Ferrario, "Renewable electricity in Europe," *Renewable and Sustainable Energy Reviews*, vol. 15, no. 8, pp. 3703–3716, Oct. 2011, doi: 10.1016/J.RSER.2011.07.015.

[5] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities," *Renewable and Sustainable Energy Reviews*, vol. 31, pp. 622–638, Mar. 2014, doi: 10.1016/J.RSER.2013.12.007.

[6] Z. Liu, M. Lin, A. Wierman, S. Low, and L. L. H. Andrew, "Greening geographical load balancing," *IEEE/ACM Transactions on Networking*, vol. 23, no. 2, pp. 657–671, Apr. 2015, doi: 10.1109/TNET.2014.2308295.

[7] V. Mathew, R. K. Sitaraman, and P. Shenoy, "Energy-aware load balancing in content delivery networks," *Proceedings - IEEE INFOCOM*, pp. 954–962, 2012, doi: 10.1109/INFCOM.2012.6195846.

[8] A. M. Omer, "Green energies and the environment," *Renewable and Sustainable Energy Reviews*, vol. 12, no. 7, pp. 1789–1821, Sep. 2008, doi: 10.1016/J.RSER.2006.05.009.

[9] Z. Cao, X. Zhou, H. Hu, Z. Wang, and Y. Wen, "Toward a Systematic Survey for Carbon Neutral Data Centers," *IEEE Communications Surveys and Tutorials*, vol. 24, no. 2, pp. 895–936, 2022, doi: 10.1109/COMST.2022.3161275.

[10] "DDoS threat report for 2023 Q2." Accessed: Nov. 22, 2023. [Online]. Available: https://blog.cloudflare.com/ddos-threat-report-2023-q2/

[11] 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, doi: 10.1109/JSAC.2016.2558918.

[12] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye, "IP geolocation databases," ACM SIGCOMM Computer Communication Review, vol. 41, no. 2, p. 53, Apr. 2011, doi: https://doi.org/10.1145/1971162.1971171.

[13] "Home - Green Web Foundation." Accessed: Nov. 24, 2024. [Online]. Available: https://www.thegreenwebfoundation.org/

[14] T. Lu, X. Lü, M. Remes, and M. Viljanen, "Investigation of air management and energy performance in a data center in Finland: Case study," *Energy Build*, vol. 43, no. 12, pp. 3360–3372, Dec. 2011, doi: 10.1016/J.ENBUILD.2011.08.034.

[15] M. Dayarathna, Y. Wen, and R. Fan, "Data center energy consumption modeling: A survey," *IEEE Communications Surveys and Tutorials*, vol. 18, no. 1, pp. 732–794, Jan. 2016, doi: 10.1109/COMST.2015.2481183.

[16] "Sustainable Innovation & Technology - Google Sustainability." Accessed: Jan. 22, 2024. [Online]. Available: https://sustainability.google/reports/google-2023-environmental-report/

[17] H. Vranken, "Sustainability of bitcoin and blockchains," Curr Opin Environ Sustain, vol. 28, pp. 1–9, Oct. 2017, doi: 10.1016/J.COSUST.2017.04.011.

[18] H. Rong, H. Zhang, S. Xiao, C. Li, and C. Hu, "Optimizing energy consumption for data centers," *Renewable and Sustainable Energy Reviews*, vol. 58, pp. 674–691, May 2016, doi: 10.1016/J.RSER.2015.12.283.

[19] W. Deng, F. Liu, H. Jin, B. Li, and D. Li, "Harnessing renewable energy in cloud datacenters: Opportunities and challenges," *IEEE Netw*, vol. 28, no. 1, pp. 48–55, Jan. 2014, doi: 10.1109/MNET.2014.6724106.

[20] K. M. Habibullah, E. Rondeau, and J. P. Georges, "Reducing energy consumption of network infrastructure using spectral approach," *Technology for Smart Futures*, pp. 235–250, Sep. 2017, doi: 10.1007/978-3-319-60137-3_11/TABLES/6.

[21] "Cloudflare radar" Accessed Jan. 4, 2024. [Online]. Available: https://radar.cloudflare.com/domains

[22] "Websiteregister Rijksoverheid | Rijkswebsites | CommunicatieRijk." Accessed: Dec. 18, 2023. [Online]. Available: https://www.communicatierijk.nl/vakkennis/rijkswebsites/verplichte-richtlijnen/websiteregister-rijksoverheid

[23] "Most Visited Websites in the Netherlands 2023 | Open .Trends." Accessed: Dec. 10, 2023. [Online]. Available: https://www.semrush.com/trending-websites/nl/all

[24] "OpenINTEL: Active DNS Measurement Project." Accessed: Nov. 14, 2023. [Online]. Available: https://www.openintel.nl/

[25] L. Universitet, +++ S.------L., F. S. Felix, and Z. Yverås, "DNS Performance : A study of free, public and popular DNS servers in 2019," 2019, Accessed: Dec. 03, 2023. [Online]. Available: https://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-158090

[26] "Website Carbon API - Website Carbon Calculator." Accessed: Dec. 22, 2023. [Online]. Available: https://www.websitecarbon.com/api/

[27] "Estimating Digital Emissions - Sustainable Web Design." Accessed: Dec. 23, 2023. [Online]. Available: https://sustainablewebdesign.org/calculating-digital-emissions/

[28] "Website Speed Test | Pingdom Tools." Accessed: Dec. 29, 2023. [Online]. Available: https://tools.pingdom.com/

[29] "Greenhouse Gas Equivalencies Calculator | US EPA." Accessed: Jan. 06, 2024. [Online]. Available: https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator

[30] M.-A. Karyotakis and N. Antonopoulos, "Web Communication: A Content Analysis of Green Hosting Companies," *Sustainability*, vol. 13, no. 2, p. 495, Jan. 2021, doi: https://doi.org/10.3390/su13020495.

[31] M. van E. Z. en Klimaat, "Strategy - Ministry of Economic Affairs and Climate Policy - Government.nl," Accessed Jan. 17, 2024. [Online]. Available: https://www.government.nl/ministries/ministry-of-economic-affairs-and-climate-policy/strategy