UNIVERSITY OF CALIFORNIA

SANTA CRUZ

IS YOUR VPN GETTING YOU BLOCKED?

A thesis submitted in partial satisfaction

of the requirements for the degree of

MASTER OF SCIENCE

in

COMPUTER SCIENCE AND ENGINEERING

by

Pakhi Sinha

August 2026

The Thesis of Pakhi Sinha
is approved:

Professor Ram Sundara Raman, Chair

Professor Chen Qian

Professor Alvaro A. Cardenas

Donald Smith

Interim Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

List of Tables

Abstract

Is Your VPN Getting You Blocked?

Pakhi Sinha

There are many valid reasons for someone to browse the internet privately, and Virtual Private Networks (VPNs) are one method of achieving this. Although primarily used to provide confidentiality and anonymity by encrypting internet traffic and masking IP addresses, users may choose to use a VPN for multiple purposes. However, this raises questions about whether, and how differently, content across the internet is portrayed to VPN users.

VPN providers do not all route their traffic the same way, and this study analyzes these differences by investigating two major VPN providers, NordVPN and Surfshark. We measure and quantify the accessibility of popular websites to users of these two VPNs, comparing outcomes against a non-VPN control connection to determine whether, and to what extent, VPN use affects a user's ability to reach and use ordinary websites.

This raises a central question: is VPN-related blocking better explained by the network a VPN happens to use, or by where its exit server is physically located? To investigate this, requests to the Tranco top 10,000 domains were compared between each VPN and a non-VPN control connection to measure blocking at full scale. A second set of experiments isolated exit servers by network (ASN) and by physical location across a smaller domain set, allowing the effects of these factors to be compared.

At full scale, NordVPN showed 1.59% VPN-attributable blocking and Surfshark

showed 0.42%. When exit servers were isolated by ASN and location, block-rate variation did not show a clear difference between the two factors, while the network path taken by a request varied substantially with location. These results suggest that VPN-related blocking is not explained by a single infrastructural factor, but instead depends on multiple properties of the network through which a VPN request is routed, including characteristics of the individual exit, network infrastructure, and geographic location.

At full scale, NordVPN showed 1.59% VPN-attributable blocking and Surfshark showed 0.42%, measured across the Tranco top 10,000 domains using a single exit per provider. Under infrastructure isolation across five ASNs and five countries, crawling top 200 domains, block-rate variation was similar whether ASN or location was held fixed (2.5% vs. 4.0% mean within-group range). Per-provider significance testing further found that only a small number of ASNs showed significant differences in blocking. In contrast, network routing varied substantially with exit location: CDN routing varied by 76% when location was held fixed, compared with 58% when ASN was held fixed. These results suggest that VPN-related blocking is not explained by a single infrastructural factor, but instead depends on multiple properties of the specific exit and network path through which a VPN request is routed.

Acknowledgments

I would like to acknowledge my advisor, Professor Ram, who has pushed me to learn and guided me throughout this process. I would also like to thank Sidharth Munigan their assistance. Finally, I am grateful to my parents who always encourage me to pursue my passions, and my friends for their unbounded support.

List of Abbreviations

VPN	Virtual Private Network
API	Application Programming Interface
CDN	Content Delivery Network
IP	Internet Protocol
ASN	Autonomous System Number
WAF	Web Application Firewall
CAPTCHA	Completely Automated Public Turing test to tell Computers and Humans Apart

Chapter 1

Introduction

Originally built for enterprise remote-access use cases, VPNs have become a mainstream consumer tool over the past decade. A VPN creates an encrypted tunnel that sits between a user's device and a remote server. Therefore, when a user visits a website, their traffic is routed through the VPN server, which forwards the request to the destination. As a result, the destination sees the VPN server's IP address rather than the user's, thus protecting the users identity [6].

Over the past decade, VPNs have become a mainstream consumer tool, and the commercial VPN market has grown accordingly. The global VPN market revenue at roughly $\$70 - 89$ billion in 2025, and is projected to have continued double-digit annual growth through the early 2030s [9, 28].

Although privacy and security remain the top reported reasons for VPN use, consumer surveys show that VPNs have evolved into a broader tool for accessing content, services, and prices that may vary by geographic location [29, 34]. Users employ VPNs to access region-blocked content, like maintain access to home-country banking or work services while traveling. Studies have shown that many e-commerce, travel, and subscription platforms price differently by

apparent geographic location– a form of consumer price discrimination– and thus switching a VPN's exit region can surface materially lower prices for flights, hotels, subscriptions [10]. Gamers similarly use VPNs to reduce latency to specific regional servers.

This breadth of use is precisely what makes investigating discrimination against VPN traffic consequential. Prior work has shown that the content and functionality a user can access depends on how their connection appears to the server [16]. This kind of blocking is often decided at the CDN level rather than by individual websites [21]. Although when content networks block VPN and anonymity traffic, the internet stops being a universally accessible resource, and what a user can see and do online varies.

This echoes a longer-running critique of the internet. With personalization algorithms heavily tailoring search results [25], VPN-targeted blocking adds a second, more blunt layer. Personalization shapes the content users are shown, while VPN blocking can determine whether that content is accessible to them in the first place. Together, these mechanisms make the internet increasingly dependent on the conditions under which a user accesses it, rather than functioning as a shared, universal resource. This is a well researched pattern called internet fragmentation, or "splintering" [23].

Fragmentation research broadly distinguishes technical, governmental, and commercial forms of this splintering [8]. Prior research largely provides insight into the behavior being observed, but there is no comparative picture yet of quantitative factors that may cause a service to block VPN traffic.

This is a difficult problem to pin down. When a website blocks a request, there may be several factors that contribute to this behavior. This is difficult to disentangle because a block may reflect several overlapping properties of a VPN

connection: geographic location of the exit IP, characteristics of its ASN, generic proxy or datacenter-IP heuristic unrelated to VPN use specifically, or a policy targeting VPN traffic. These elements are typically confounded in a single-provider study design, and thus, it is difficult to isolate the factors that influence blocking.

This thesis asks a simple question: when a website blocks a VPN user, is it blocking the VPN, or the network the VPN happens to be routed through? I work toward closing that gap by empirically comparing website accessibility across two major commercial VPN providers, NordVPN [24] and Surfshark [32], using a paired measurement approach. I crawl a ranked, stability-hardened list of popular websites from the Tranco list [19], requesting each domain both through a VPN connection and without one. The two outcomes are compared to determine whether the VPN connection was blocked, degraded, or unaffected, adapting a block-classification approach previously used to measure VPN accessibility [3]. I focus on analyzing two main factors– an exit's ASN versus its geographical location. To isolate these factors, I sample exit servers in two complementary ways: one holds the ASN fixed while varying the location, while the other holds the location fixed while varying the ASN. This allows me to examine whether blocking is associated with where an exit is located or which network it belongs to, rather than attributing the behavior to the VPN provider as a whole. Each request's exit IP, ASN, and location are recorded alongside its access outcome, allowing blocking behavior to be examined in terms of the underlying network infrastructure.

At full scale, both providers showed measurable VPN-attributable blocking, though the rate differed considerably between them. When exit servers were further isolated by ASN and by location across a smaller domain set, block-rate variation was comparable regardless of which factor was held fixed, suggesting that no single infrastructural property cleanly accounts for blocking on its own.

However, the network path a request was routed through- specifically, which content delivery edge it reached- varied far more with exit location than with ASN. This indicates that ASN and location play distinct, sometimes opposing roles depending on which aspect of VPN-mediated web behavior is being measured, rather than one factor dominating uniformly.

This thesis makes the following contributions:

- A comparative empirical measurement of website blocking across two major commercial VPN providers, extending a prior single-provider methodology to a multi-provider, infrastructure-aware design.
- An experimental design that isolates a VPN exit's ASN and its geographic location as distinct explanatory factors, rather than treating VPN blocking as a property of the provider as a whole.
- Empirical evidence that competing VPN providers frequently share underlying hosting infrastructure, complicating the assumption that different providers constitute independent vantage points.
- A discussion of the practical implications of these findings for VPN providers' exit-server selection, CDN and WAF blocking policy design, and end users navigating VPN-related access issues.

Overall, this work aims to move the study of VPN-related web discrimination from a per-provider accounting of blocking rates toward an infrastructure-aware account of why blocking happens. This distinction matters because VPNs are fundamentally dual-use: the same infrastructure that lets an ordinary user protect their privacy also lets a malicious actor hide their identity. Current blocking

heuristics— largely IP and ASN-based— cannot easily tell these two cases apart, which is part of why they end up penalizing genuine users along with bad actors.

Addressing this likely requires two things beyond what this thesis measures directly: longer-term, continuous monitoring, since blocking behavior can change as ASNs are re-flagged or reputations shift over time, and progress toward more standardized, less ad hoc access policies across CDNs, so that a user's experience does not depend arbitrarily on which network happens to host their exit IP. A more durable fix likely lies beyond IP-based blocking altogether. Privacy Pass [7] allow a server to distinguish trustworthy from abusive traffic without relying on the reputation of the underlying IP or ASN at all by having a client earn a cryptographic token after passing a challenge once. This token is then redeemed on future requests. Although such approaches may have limitations, they point toward a future where genuine VPN users are not collaterally penalized for the infrastructure they happen to route through.

Chapter 2

Background and Related Work

This chapter reviews existing literature on VPNs and anonymity-network blocking, connects that work to the broader problem of internet fragmentation. It examines how prior studies have measured VPN accessibility and characterized the infrastructure underlying web-based blocking, while identifying the limitations of existing approaches. Finally, it situates this thesis within that literature and motivates the need for a comparative, infrastructure-aware analysis of the factors associated with VPN blocking.

2.1 Privacy, Anonymity Networks, and Server-Side Discrimination

The right to online privacy has been recognized as a human right by the United Nations [4]. This matters because it frames tools such as Tor and commercial VPNs not simply as technical utilities, but as ways for users to exercise this right in practice. This is especially important for journalists, activists, and users living under censorship, for whom anonymity can be necessary for safe participation

online [13, 12]. At the same time, the properties that protect a user's privacy—a shared IP address and traffic routed through a third party—can also make that traffic difficult for websites to distinguish from abusive traffic.

A growing body of work shows that websites and content delivery networks (CDNs) resolve this ambiguity by treating anonymized IP addresses as inherently suspect, applying blocking, CAPTCHAs, or degraded content uniformly rather than differentiating genuine users from malicious ones [7]. In other words, the privacy protection a VPN or anonymity network offers and the accessibility penalty it imposes are two sides of the same infrastructural coin: a user cannot easily have one without risking the other, and the scale of that risk is precisely what this thesis sets out to measure.

2.2 VPN Security, Privacy, and Performance Research

Before asking how websites treat VPN traffic, it is worth examining a quieter assumption behind that question: that a VPN reliably provides the protection it promises. Existing research suggests that this cannot be taken for granted. Across security, privacy, and performance, studies have found substantial differences between VPN providers and their underlying implementations. This matters for the present thesis because it suggests that "using a VPN" does not describe a single, uniform technical condition. Instead, the behavior and risks of a VPN connection depend on the specific provider, client, protocol, and infrastructure being used.

2.2.1 Security

Security research has repeatedly found gaps between what users expect a VPN to do and what VPN clients actually do. Ikram et al. [11] analyzed 283 Android applications that request VPN permissions and found that many contained third-party tracking libraries, malware, or mechanisms that manipulated user traffic. Perta et al. [27] identified another important failure: IPv6 traffic leakage and DNS hijacking across popular commercial VPN clients. These leaks can be difficult for users to detect because the VPN connection can appear active while some traffic bypasses the encrypted tunnel. Khan et al. [15] found similar differences across commercial VPN applications, showing that these problems are not limited to a single provider or implementation.

Together, these findings show that VPN security is not a fixed property of using a VPN. It depends on how a particular provider implements its client and tunnel. Two users may both believe they are protected by a VPN while experiencing different levels of protection depending on the application they use.

2.2.2 Privacy

Privacy research raises a related question: what does a VPN provider do with the traffic it sees? Khan et al. [15] found substantial differences between commercial VPN providers in areas such as logging, DNS handling, encryption practices, and third-party tracking. They also highlight that many commercial VPN applications embed third-party trackers directly within the client, meaning a portion of user data is exposed not to the open web, but to the VPN provider's own advertising or analytics partners, an arrangement invisible to a user who chose the VPN specifically to avoid such tracking.

VPNalyzer [17] recognizes that commercial VPN providers differ substantially in network configuration, logging practices, and privacy posture, despite being marketed and perceived by consumers as an umbrella privacy tool. These differences include the ASNs and hosting arrangements used by providers. These findings suggest that a user's choice of provider is just as important while determining how much privacy they actually receive.

2.2.3 Performance

Performance shows a similar pattern of provider-dependent variation. Early comparative work on software VPNs found that encryption overhead can degrade throughput substantially and increase CPU load. The severity of the impact depends on the specific protocol and cipher configuration used rather than on any property inherent to VPN use itself [26]. More recent comparisons of VPN protocols, including OpenVPN and WireGuard, continue to find differences in throughput, latency, and resource consumption depending on the protocol and deployment environment [18].

Taken together, this literature converges on a single point directly relevant to this thesis: security, privacy, and performance research all independently conclude that commercial VPN providers are not interchangeable, and that provider-specific implementation and infrastructure choices determine outcomes. This matters for how blocking research should be designed. If VPN providers are already known to diverge this sharply across every other dimension that has been measured, a study that treats a single provider's blocking rate as representative of VPN users broadly is making an assumption the rest of the VPN literature has already shown to be unreliable. The remainder of this chapter examines how prior blocking research

has nonetheless proceeded largely on a per-provider basis, and what is lost by doing so.

2.3 Measuring Blocking of Anonymity Network Users

Most of what we know about anonymity-network blocking comes from studying Tor, not VPNs, which matters because Tor and VPNs are structurally different. Tor exit IPs are public and well documented, while VPN exit IPs are commercial infrastructure, often quietly shared with unrelated hosting customers.

Khattak et al. [16] were among the first to systematically measure this for Tor, showing that popular websites treat Tor exit traffic differently through two main mechanisms: outright HTTP-level blocks, and more subtle content-level discrimination, where the site loads but shows something different (a CAPTCHA, a stripped-down page, or a warning). This distinction matters because it means "blocked" is not a single, easy-to-detect outcome — a user can technically reach a site and still not really be able to use it, which is exactly the kind of subtlety a blocking study needs to capture rather than assume away.

Singh et al. [30] took this further by examining why Tor exits are blocked. They identified two mechanisms that can produce the same outcome: reactive blacklisting, where a shared IP is blocked after abusive behavior, and proactive blacklisting, where an IP is blocked simply because it is a known exit node. This distinction is directly relevant to VPNs, where shared, datacenter-hosted exit IPs can be subject to the same mechanisms. A VPN user may therefore be blocked not because of their own behavior, but because their exit IP was previously flagged or

is treated as suspicious by default. Distinguishing between these mechanisms is important because it determines what is actually driving the block.

Direct VPN blocking measurement is much rarer. The closest and most direct precedent for this thesis is a series of bachelor's theses from TU Delft [3], each applying the same basic method— request a website with and without an anonymity tool, then compare the results. They found no meaningful blocking when only requesting a website's homepage (0.97% of about 15,000 requests), but a significant jump once subpages were also checked (1.38% of domains). Blocking was especially common on essential services like health and government sites. This result matters practically: a study that only checks homepages would have concluded VPNs are barely blocked at all, missing a real and meaningful problem that only shows up once you look slightly deeper. It is a strong argument for why this thesis also tests beyond the homepage.

A common limitation across prior research here is that VPNs and anonymity networks are often treated as single, uniform entities. A block observed while testing a provider is typically attributed to that provider, without examining whether the outcome is instead driven by properties of the underlying infrastructure, such as the hosting network, ASN, or IP address. A block found while testing Mullvad is reported as "Mullvad was blocked," and does not answer whether Mullvad specifically was blocked, or the hosting company whose IP address Mullvad happened to be using at the time. This thesis addresses this gap by separating provider identity from these underlying network characteristics.

2.4 CDN Geoblocking, IP Reputation, and Infrastructure-Level Bias

McDonald et al. [21] shows that a large share of geoblocking on the modern web is not decided by individual websites at all, but by the content delivery networks (CDNs) that sit in front of them. This shows that one CDN's blocking policy can quietly affect thousands of unrelated websites at once.

Tschantz et al. [33] broaden this picture further, cataloging the many different reasons a site might become unreachable, including automated bot-detection systems that specifically distrust datacenter-hosted IP ranges. This reinforces that from the server's perspective, a VPN exit may therefore be indistinguishable from other suspicious datacenter traffic, regardless of the behavior of the individual user.

More recent works explore this infrastructure-level view, rather than a VPN-identity view. Master and Garman [20] document a shift away from blanket IP-based blocking, in part because blocking an entire shared CDN IP range can cause significant collateral damage to legitimate users. This suggests that blocking systems are becoming more aware of the specific IP and ASN associated with a request. Khan et al. [14] ("Stranger VPNs") find a similar pattern: commercial VPN providers differ substantially in their ability to unblock geographically restricted content, and these differences are linked to two main infrastructural causes– a meaningful share of exit servers are physically hosted in a different country than the one advertised to users, and many exit IPs are independently identifiable and pre-flagged as known VPN or proxy addresses by streaming services' own detection systems, regardless of server location. The second cause in particular mirrors the exact mechanism this thesis investigates for blocking generally: an IP's reputation

as VPN infrastructure, not the provider's identity, appears to drive the outcome.

A more recent study [2] shows that even when two providers offer an exit in the same country, they may route traffic through different servers, hosts, and physical locations, producing different observations of the web, affecting the measurements researchers obtain.

2.5 Summary

Pulling these threads together: blocking research knows a great deal about whether anonymity tools get blocked, but almost always studies one tool at a time, without asking whether the network underneath that tool is what actually matters (§2.3). Separately, infrastructure and CDN research knows a great deal about why IP-based blocking happens in general, but rarely connects that back to VPN blocking specifically (§2.4). And a growing body of work shows that VPN providers are not interchangeable to begin with, whether in security, privacy, performance, or even the basic network path a measurement takes (§2.2, §2.5). No existing study sits at the intersection of all three: measuring VPN blocking specifically, across more than one provider, while treating the underlying ASN and IP infrastructure as the variable that actually explains the outcome.

This thesis is positioned directly in that gap. By comparing NordVPN and Surfshark using a shared measurement methodology, and by deliberately sampling exit servers to separate the effect of network (ASN) from the effect of location, this work aims to move VPN-blocking research from a per-provider accounting of block rates toward an infrastructure-aware explanation of why those blocks happen in the first place.

Chapter 3

Methodology

3.1 Data Collection Pipelines

Data for this study was collected across three pipelines, run in sequence as the experimental design matured. The first two pipelines are single-exit crawls: one without a VPN (the control) and one through a single VPN server per provider, each run against the full Tranco top-10,000 domain list. The third pipeline is a smaller, more tightly controlled crawl designed to separate two variables that are normally confounded in VPN measurement: the exit server's ASN (which hosting company owns the IP) and its geographic location.

3.1.1 Pipeline A- Control Baseline

This pipeline crawls all 10,000 domains with no VPN active. It is the baseline everything else is compared against.

Pipeline A establishes the control condition by crawling all 10,000 domains without an active VPN connection. The same crawler and measurement procedure used in Pipelines B and C are applied to the control crawl, providing a matched

baseline against which VPN observations can be compared. For each domain, the crawler records the measurements described in Table 3.1.

3.1.2 Pipeline B- Full-Scale Single-Exit Crawl

This pipeline asks the simplest version of the research question: does connecting through a VPN, from one ordinary exit server, change what a website shows or does, across a large number of sites?

Both NordVPN and Surfshark were connected once, to a single exit each, and left connected for the full 10,000-domain crawl. No automatic exit-switching happens here.

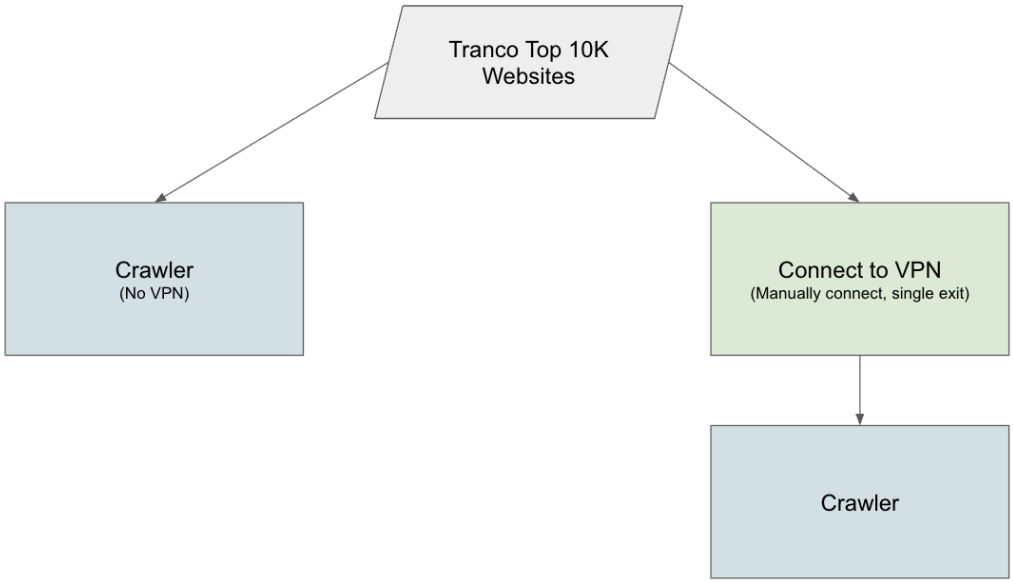

Figure 3.1: The manual, single-exit pipeline.

3.1.3 Pipeline C- Isolating ASN from Location

A single VPN exit is not representative of an entire provider. Exits span multiple countries and hosting companies, making it difficult to determine whether blocking

is associated with location, hosting provider, or both. Pipeline C separates these factors by varying one while holding the other fixed:

- **Fix hosting company, vary location.** The five largest hosting companies are selected, with five exits per company chosen across different countries.
- **Fix location, vary hosting company.** The five countries containing the most distinct hosting companies are selected, with one exit per hosting company within each country.

Each approach yields 25 exits per provider. To keep the crawl feasible, Pipeline C tests 200 domains per exit rather than all 10,000. Thus, Pipeline B provides broader domain coverage but cannot disentangle hosting company from location, while Pipeline C separates these factors at the cost of reduced domain coverage.

Both providers are evaluated using the same crawler in Pipeline C, including detection of hard blocks and interactive challenges such as CAPTCHAs. This enables direct comparison between NordVPN and Surfshark within this pipeline.

3.2 Implementation and Provider-Specific Constraints

NordVPN and Surfshark differ in how precisely individual exit servers can be selected, which affects how Pipeline C is executed. NordVPN's connection tool supports direct connections to named servers, allowing each sampled exit to be selected explicitly by IP and hosting company as shown in Figure 3.2a.

Surfshark's client does not provide equivalent server-level selection; it only supports connecting to the nearest server. However, individual exits have corresponding OpenVPN configuration files. Pipeline C therefore bypasses the Surfshark

client and connects directly through these configurations. This required additional handling to maintain administrator access, preserve the remote SSH connection when VPN routes were established, and verify connection success from OpenVPN logs.

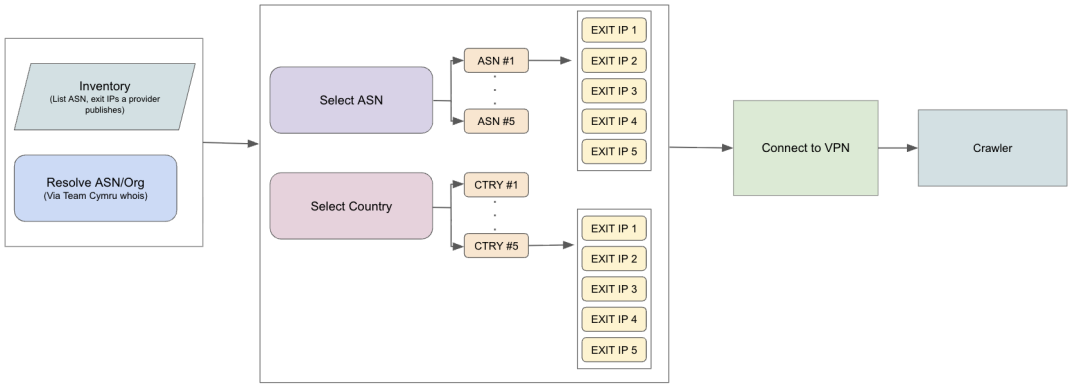

(a) NordVPN's infrastructure-isolation pipeline. Exits are inventoried, sampled by ASN and location, then connected to directly by server name before comparison.

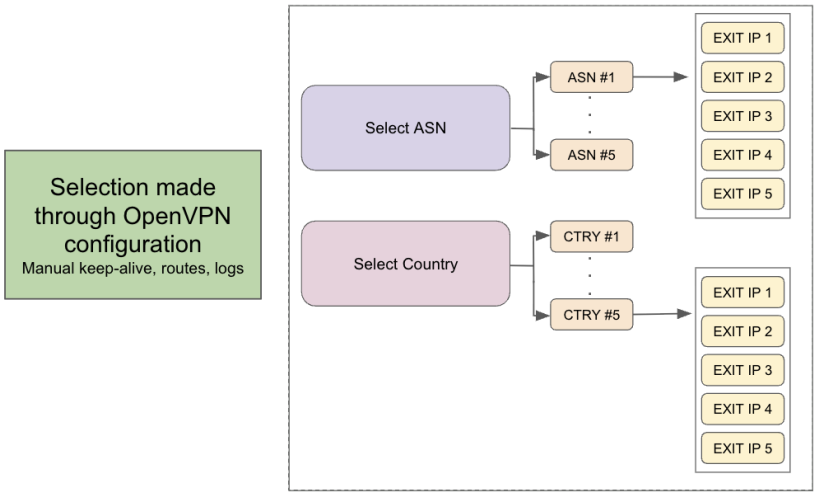

(b) Surfshark's infrastructure-isolation pipeline. Differs only in how selection is made.

Figure 3.2: Infrastructure-isolation pipelines for both providers.

The direct OpenVPN approach required several additional operational safeguards:

- **Privilege persistence.** Administrator access was refreshed automatically during long multi-exit runs to prevent the crawl from stalling on password prompts.
- **SSH route preservation.** Because establishing a VPN tunnel can replace the machine's default route, the existing SSH route was identified and protected before each connection to prevent the VPN from disconnecting the remote crawl process.
- **Connection verification.** Unlike the NordVPN client, OpenVPN does not provide an equivalent connection-status interface. Connection success was therefore determined by monitoring OpenVPN logs for known success and failure messages.

These differences are operational rather than experimental: they affect how reliably selected exits can be connected to, but do not constitute differences in the blocking measurements themselves. They are therefore treated as a methodological limitation (Section 5.4). Both pipelines also support automatic resumption, skipping previously completed domains and exits after an interruption.

3.3 Metrics and Data Schema

The pipeline records data at three levels of granularity, corresponding to the three things that need to be known to answer the research question: what happened on a given domain, which exit produced that result, and how that result differed from

the no-VPN baseline. Table 3.1 covers the first, Table 3.2 the second, and Table 3.3 the third. Together, these three layers let a single blocked request be traced back to the exit that produced it and weighed against what the same domain looked like without a VPN.

3.3.1 Per-Domain Record

Every domain visited produces one record. Each record includes:

Measurement	**What We Record**
Domain and navigation	Domain identity, Tranco rank, requested and final URLs, browser, and crawl timestamp.
Page loading	Total page load time, time to first byte, navigation timeout, network-idle status, and HTTP status.
Network activity	Total requests, third-party requests, probable tracker requests, and failed resources. Measurements are recorded separately for the homepage and the complete visit.
Cookies	Cookies observed before and after post-load actions, including first- and third-party cookies.
Visual appearance	Perceptual hashes of full-page screenshots captured at homepage load and after post-load actions.
User interactions	Actions performed after the initial page load, including scrolling, cookie-consent interaction, and navigation to one safe internal page.
Blocking indicators	For potentially blocked responses, page markers, challenge indicators, HTTP status, and detected security or content-delivery providers. A raw response snippet is also retained for inspection.
Network identity	Observed exit IP, organization, and ASN of the crawling machine.

Table 3.1: Measurements Collected During Each Domain Crawl

3.3.2 VPN Exit Metadata

Each VPN record is also linked to its specific exit, allowing observed blocking patterns to be analyzed by ASN and location rather than provider alone. In Pipeline C, mismatches between the intended and observed ASN are flagged to ensure that comparisons use the sampled exit as intended.

Measurement	**What We Record**
Exit server	VPN server hostname and selected exit information.
Exit location	Provider-reported country, WHOIS country, and continent.
Network identity	Observed exit IP, organization, and ASN after establishing the VPN connection.
Connection outcome	Whether the VPN connection and subsequent crawl completed successfully.
Exit validation	Whether the observed ASN matched the ASN associated with the selected exit.

Table 3.2: VPN Exit Metadata

3.3.3 Analysis Metrics

The final layer merges control and VPN records for each domain into a single comparison. The study's conclusions are based on differences between these conditions, rather than either condition alone.

Metric	Description
Availability	Whether the domain was successfully reached in the control and VPN conditions.
Load-time difference	Difference in page load time between VPN and control, measured as VPN minus control.
TTFB difference	Difference in time to first byte between VPN and control, measured as VPN minus control.
Redirect difference	Difference in the number of navigation redirects between VPN and control.
Failed-resource difference	Difference in failed resource count between VPN and control.
Screenshot distance	Perceptual difference between VPN and control screenshots, based on their perceptual hashes.

Table 3.3: VPN–Control Comparison Metrics

3.4 Exit Inventory and Target Selection

Before running the multi-exit crawl, an inventory is constructed for each provider's available exit servers. Each entry records the exit IP address, provider-reported country, country returned by an independent registry lookup, hosting organization, and corresponding ASN. These fields are kept separate because the country advertised by a VPN provider and the location associated with the underlying IP address may not be the same.

Rather than selecting ASNs at random, the most populated ASNs within each provider's exit inventory are identified and used as the basis for target selection. For NordVPN, the server inventory is obtained through its public server API, which provides information about available servers and their locations [1]. For Surfshark, the inventory is constructed from the available OpenVPN configuration files, which provide the provider's server endpoints and locations. The resulting inventories are grouped by ASN, and the most represented ASNs for each provider

are selected for further measurement. This focuses the analysis on networks that account for a substantial share of each provider's available exits.

The selected exits are then sampled according to the approach used in Pipeline C (Section 3.1.3). In one sampling scheme, the ASN is held fixed while the location is varied; in the other, the location is held fixed while the ASN is varied. This allows the two factors to be examined separately rather than treating the VPN provider as the unit of analysis.

The inventories also show that the same underlying network infrastructure can appear across providers. The same ASN and hosting organization may therefore occur in both the NordVPN and Surfshark inventories. This is not an error: an ASN identifies the network that announces an IP address, rather than the VPN provider using it. For example, AS212238 is associated with CDNEXT/Datacamp Limited, while AS9009 is associated with M247. When both providers use exits within these networks, the same ASN and organization appear in both inventories. The providers may nevertheless use different IP addresses, server endpoints, locations, and subsets of the available IP space. This overlap is relevant to the experimental design because it demonstrates why VPN provider identity cannot be used as a direct proxy for underlying network infrastructure.

3.4.1 Sweep Record

For each exit attempted in Pipeline C, the following information is recorded:

- **Connection status:** whether the VPN connection succeeds.
- **Exit information:** the observed exit IP, ASN, hosting organization, and location.

- **Target comparison:** whether the observed ASN and hosting organization match the values recorded during target selection.

Any mismatch is flagged for review, and the exit can be excluded or retried. This check ensures that the network being measured matches the network selected by the sampling procedure, which is necessary for Pipeline C to isolate the intended infrastructure-level factors.

3.4.2 Comparing Control and VPN Runs

The final analysis pairs the control and VPN measurements for the same domains. Each resulting row contains:

- whether the domain was successfully recorded under each condition;
- differences in load time, time to first byte, redirects, and failed requests;
- differences in page content and user experience, including screenshot similarity, cookie-consent banners, and third-party tracking cookies.

These paired measurements are used to summarize how often outcomes change between the control and VPN conditions, including cases where a page that succeeds under the control becomes inaccessible through the VPN. Network errors, load-time differences, and content or user-experience changes are also examined to identify potential degradation or blocking.

These measurements are treated as evidence consistent with blocking rather than definitive proof of it, since normal website variability and unrelated network failures can produce similar outcomes (see Section 5.4).

Chapter 4

Results

4.1 Overview of Findings

Across all three experimental pipelines, the results showed that VPN traffic did not receive uniform treatment across websites or VPN exits. Blocking varied between providers and individual exits, while exit ASN and location were also associated with differences in network routing and user-facing content. The following sections examine these results in detail for each experimental pipeline.

4.2 Blocking at Full Scale

This section presents the results of the full-scale crawl of the Tranco top 10,000 websites. Each domain was tested under three conditions: a non-VPN control, one NordVPN exit, and one Surfshark exit. The same domain set was used across all three conditions, allowing the VPN crawls to be compared directly against the control.

4.2.1 Access Failures and HTTP Blocking

The two VPN exits produced substantially different numbers of access failures across the 10,000 domains. As shown in Figure 4.1, the NordVPN exit produced 176 VPN-exclusive access failures and 111 cases in which a domain returned an HTTP 403 response, compared with 60 VPN-exclusive access failures and 33 HTTP 403 responses for the Surfshark exit. Thus, the NordVPN exit produced approximately two to three times as many VPN-specific failures as the Surfshark exit.

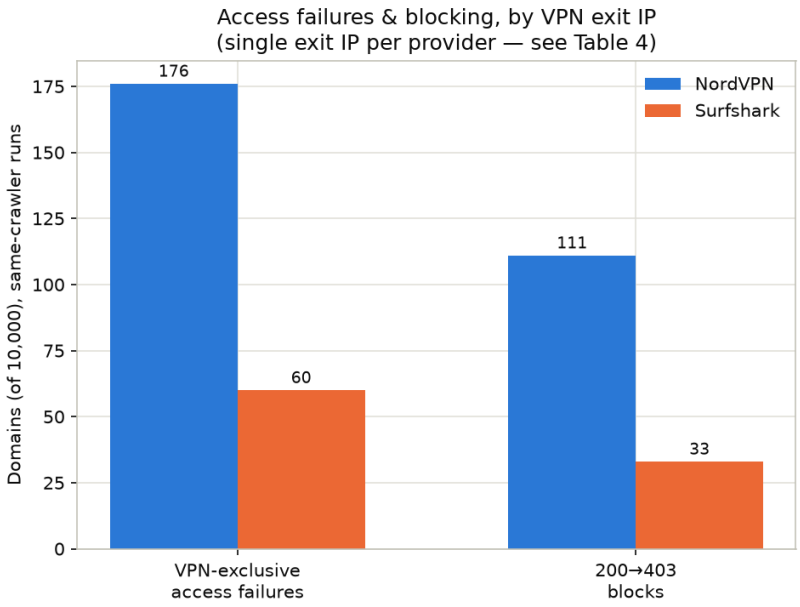

Figure 4.1: VPN-exclusive access failures and HTTP 403 blocks across the 10,000-domain crawl for the selected NordVPN and Surfshark exits.

These results indicate a substantial difference in website accessibility between the two tested exits. However, the observed difference reflects the behavior of these particular exit networks and should not be interpreted as a provider-wide difference. At this scale, the result is more appropriately viewed as evidence of differing behavior associated with the two tested exit infrastructures.

4.2.2 Blocked Domains by Sector

The HTTP 403 responses were further examined by grouping the affected domains into five broad categories:

1. Finance/Commerce
2. Government/Institutional
3. Media/News
4. Social/Platform
5. Other/Niche

Figure 4.2 shows the resulting distribution for the two VPN exits. Categories are manually assigned and are intended as qualitative groupings rather than a formal taxonomy.

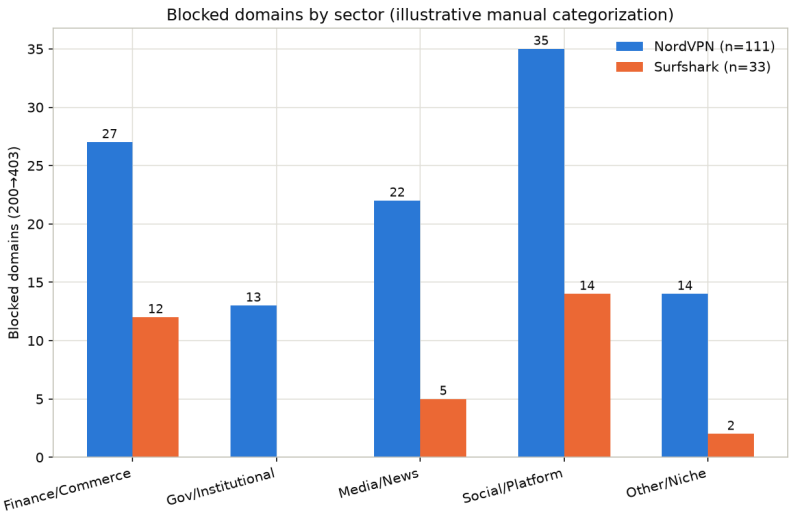

Figure 4.2: HTTP 403-blocked domains categorized by broad website sector for the NordVPN and Surfshark exits.

The distribution differs between the providers, with Social/Platform and Finance/Commerce representing the largest groups among the domains returning

403 responses for the NordVPN exit. The Surfshark exit produced substantially fewer 403 responses overall, resulting in smaller counts across the categories.

These categories are intended as a qualitative characterization of the affected domains rather than a formal website taxonomy. They were assigned manually using a set of category overrides and should therefore be interpreted as broad sector-level groupings rather than precise estimates of blocking rates by industry.

4.2.3 Overlap in Blocking Across VPN Providers

The domains blocked through the two VPN exits show limited overlap. As shown in Figure 4.3, 95 domains were blocked only through the NordVPN exit, 17 were blocked only through the Surfshark exit, and 16 were blocked through both. Thus, only 16 of the approximately 128 unique blocked domains were common to both VPN conditions.

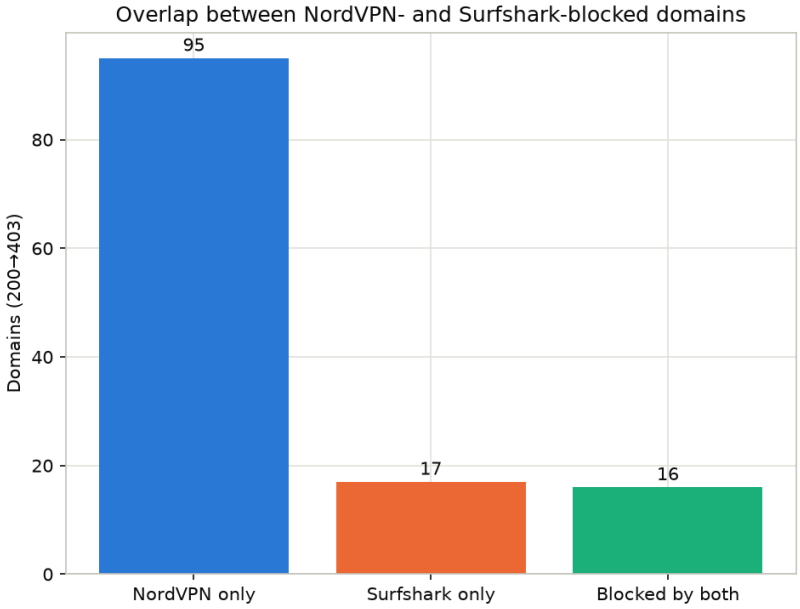

Figure 4.3: Overlap in HTTP 403-blocked domains between the selected NordVPN and Surfshark exits. Most blocked domains are unique to one exit, with only 16 domains blocked through both.

This observation suggests that website blocking is not determined solely by whether traffic originates from a VPN, since the two VPN exits produced substantially different blocking patterns. However, the experiment cannot determine which characteristic of the exits accounts for this difference. The exits differed in their IP addresses, autonomous systems, geographic locations, and potentially other network properties.

The full-scale crawl therefore establishes an important baseline: blocking varies substantially between individual VPN exits, both in overall volume and in the specific domains affected. The limited overlap between the two exits provides motivation for examining whether this variation is associated with characteristics of the underlying VPN infrastructure. In particular, the following experiments isolate ASN and geographic location to determine whether either factor can account for the observed differences in blocking.

4.2.4 Comparison to the Mullvad Precedent

Biazin do Nascimento [3] used Mullvad VPN and a smaller domain set (approximately 1,000–3,000 domains per stage), with separate measurements of homepages and subpages. They reported a homepage blocking rate of 0.97% of 15,000 requests, increasing to 1.38% when subpage access was included. In comparison, the present study found 1.59% VPN-attributable blocking for NordVPN and 0.42% for Surfshark. Thus, NordVPN was roughly 1.6× higher and Surfshark 2.3× lower than the earlier result. Overall, the observed rates span the range reported by Biazin do Nascimento, but the differences should not be interpreted as evidence of changes in VPN blocking over time, since differences in domains, methodology, VPN infrastructure, and measurement period may explain the variation.

4.3 Blocking Under Infrastructure Isolation

The infrastructure-isolation experiments extended the full-scale crawl by examining a smaller, controlled sample of 200 domains across multiple VPN exits. For each provider, 25 exits were selected across multiple autonomous systems and geographic locations. This reduced domain coverage relative to the 10,000-domain crawl, but allowed the analysis to examine variation across individual VPN exits rather than comparing only one exit per provider. The experiments therefore provide a more direct test of whether observed blocking differences are associated with ASN or geographic location.

4.3.1 Exit Infrastructure Reliability

The observed exit infrastructure did not always correspond exactly to the infrastructure specified during sampling 4.4. In particular, some connections exited through an ASN different from the one associated with the selected server. This occurred more frequently for Surfshark than NordVPN, indicating that the advertised server infrastructure cannot always be assumed to correspond to the network observed at the destination.

For NordVPN, mismatches occurred primarily for a small number of selected ASNs, while Surfshark mismatches were distributed across several ASNs.

These observations demonstrate that the selected or advertised ASN cannot always be treated as equivalent to the network from which the crawler actually exits. Consequently, subsequent blocking analyses use the observed exit infrastructure when evaluating the relationship between ASN and website behavior.

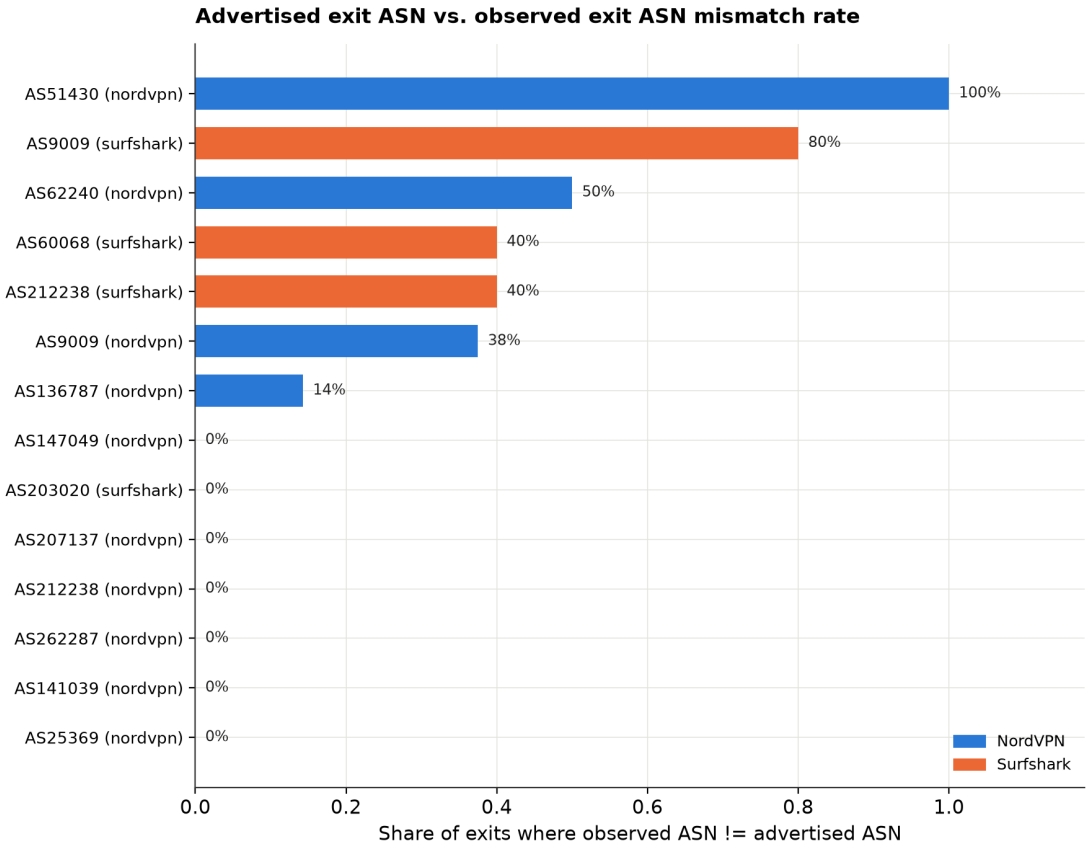

Figure 4.4: Observed ASN mismatches across the infrastructure-isolation experiments. The selected ASN is compared with the ASN reported for the IP address observed at the crawler.

4.3.2 Variation in VPN-Attributable Blocking Across ASNs

The no-VPN control provides a baseline for interpreting blocking in the infrastructure-isolation experiments. Of the 200-domain sample, 12 domains returned HTTP 403 responses in the control crawl. These domains were excluded when calculating VPN-attributable blocking, since their failures were not specific to VPN use.

After excluding control-blocked domains, VPN-attributable block rates varied across the observed ASNs, ranging from 4.3% to 8.2%. Figure 4.5 shows the distribution of block rates across the sampled infrastructure for both providers.

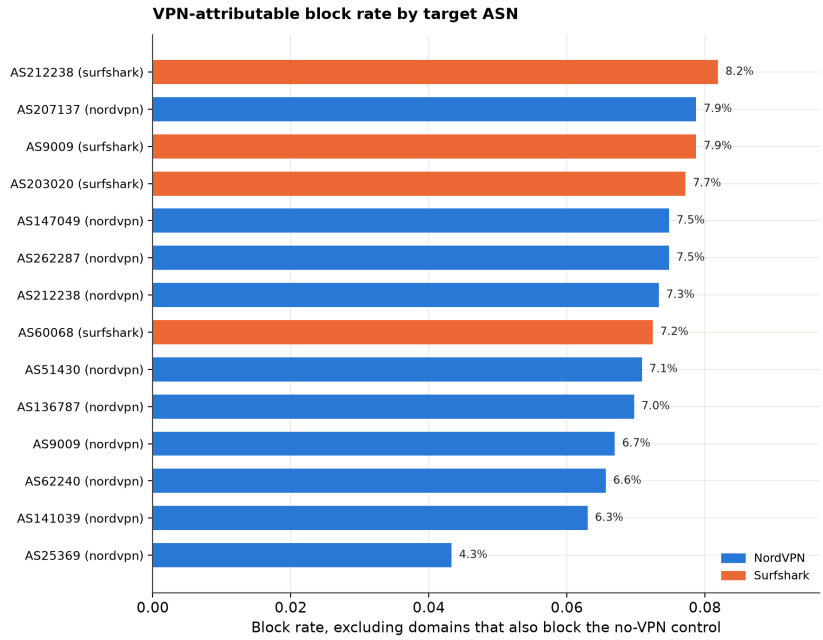

Figure 4.5: VPN-attributable block rate by observed ASN for the sampled NordVPN and Surfshark exits. Rates are calculated after excluding domains that returned HTTP 403 responses in the no-VPN control.

The observed variation is relatively modest, with most ASN-level block rates falling within a narrow range. In addition, the number of exits represented by

each ASN varies substantially, with several ASNs represented by only a single exit. The complete breakdown of sampled networks and exit counts is provided in Appendix 7.2.

Although one ASN exhibited the highest observed block rate, this descriptive difference alone does not establish an ASN-specific blocking effect. Statistical comparisons between each ASN and the remaining sampled exits found only one statistically significant difference. This result was based on four exits and should therefore be treated as suggestive rather than conclusive.

Overall, the results provide limited evidence for an ASN-specific blocking effect. While blocking rates vary across ASNs, the available sample does not establish that ASN is a dominant determinant of whether a website blocks VPN traffic.

4.3.3 ASN Effect on Blocking

The infrastructure-isolation experiments provide limited support for an ASN-specific blocking effect. Blocking rates varied across the sampled infrastructure. The results therefore do not support a simple model in which ASN reputation is the primary determinant of VPN blocking. Instead, it suggests that blocking may depend on a combination of factors associated with the exit, including its IP address, network infrastructure, geographic location, and site-specific policies. The experiment cannot isolate these factors completely. Figure 4.6 summarizes the ASN-level results across three measurements: homepage blocking, statistical comparisons of block rates, and subpage blocking.

Homepage and subpage block rates show limited and inconsistent evidence of an ASN-specific blocking effect, while CDN edge matching demonstrates that ASN is associated with other aspects of network behavior. The statistical comparisons

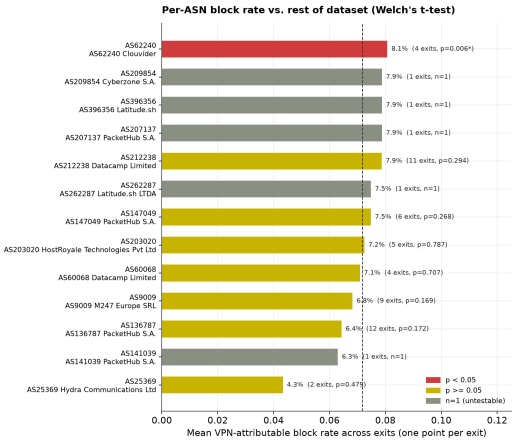

(a) Homepage blocking

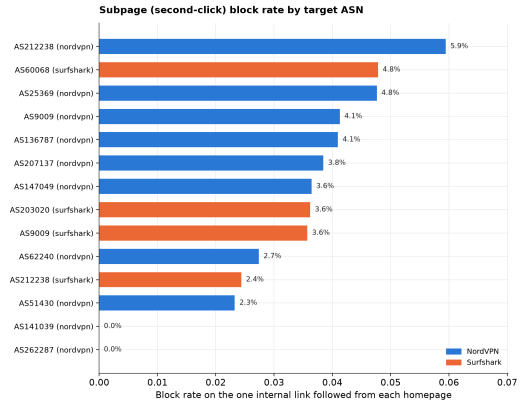

(b) Subpage blocking

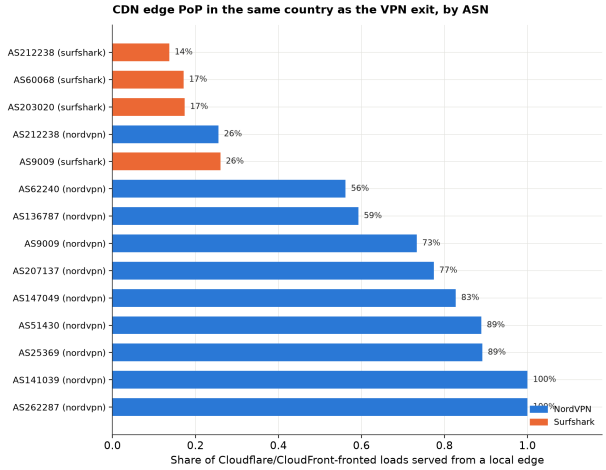

(c) CDN edge matching

Figure 4.6: ASN-level variation across blocking and network-path measurements.

provide little evidence that the observed differences in blocking rates are systematic. Only one ASN-level comparison reached conventional statistical significance, and this result was based on a small number of exits. The remaining comparisons did not distinguish their respective ASN from the rest of the sampled infrastructure.

The subpage measurements provide a similar result. Although the relative block rates varied across ASNs, the differences remained modest and did not provide strong evidence of a consistent ASN-level effect. This is notable because the subpage crawl provides a second opportunity to observe blocking beyond the initial homepage request.

The CDN measurements provide additional context for interpreting these results. While they do not measure blocking directly, CDN edge matching varied substantially across the sampled ASNs, indicating that ASN can be associated with differences in how traffic is routed through network infrastructure. This demonstrates that the lack of a strong ASN effect on blocking does not imply that ASN is irrelevant to other aspects of network behavior.

Homepage and Subpage Blocking

To determine whether the observed blocking patterns persisted beyond the initial request, the crawler followed one on-site link from each successfully loaded homepage. Across the infrastructure-isolation experiments, 2,513 subpage requests were attempted. The resulting subpage block rate was 3.8%, while 12.4% of requests failed without receiving a response.

Subpage blocking was substantially lower than homepage blocking in the sampled data. The relative ordering of several ASNs was broadly similar between the homepage and subpage measurements, see Figure 4.6b. Although, the smaller number of observations for individual exits limits the strength of this comparison.

Because subpage requests were only attempted after a homepage successfully loaded, their blocking rate is not directly comparable to the homepage block rate. The subpage results therefore provide a secondary observation rather than an independent estimate of website blocking.

4.3.4 Geographic Effects on Web Behavior

The infrastructure-isolation crawl also identified several effects associated with the geographic location of the VPN exit. Among successfully loaded pages, EU-/EEA/UK/CH exits produced cookie-consent banners on 7.2% of loads, compared with 3.6% for exits elsewhere. They also resulted in fewer third-party tracking cookies on average (5.0 versus 8.4) as shown in Figure 4.7.

Geographic steering was also observed for a small number of websites. In particular, Microsoft properties served different language or regional variants depending on the exit country, including differences in locale codes such as `en-gb` and `fr-be`. However, broad geographic redirection was not observed consistently across the websites tested.

Finally, exit location did not necessarily correspond to the location of the CDN infrastructure serving the response. Across Cloudflare and CloudFront-served requests, only 48.1% were handled by an edge in the same country as the VPN exit. This correspondence also varied across the sampled network infrastructure, as shown in Figure 4.6c.

These results show that VPN exit location can affect website presentation, tracking behavior, and network routing in addition to accessibility. Because these geographic effects are secondary to the study's primary blocking analysis, they are not examined further here.

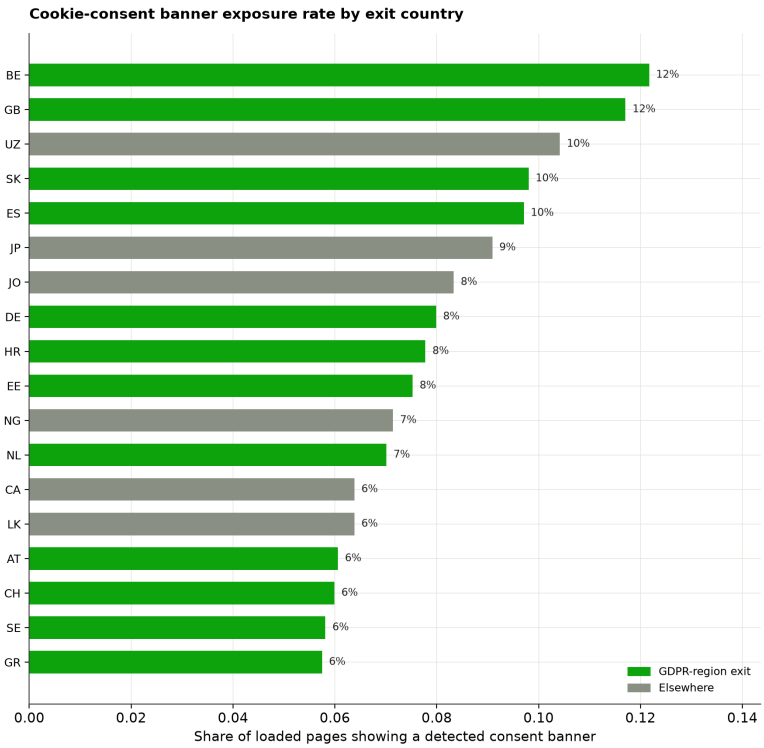

(a) Cookie-consent banner frequency

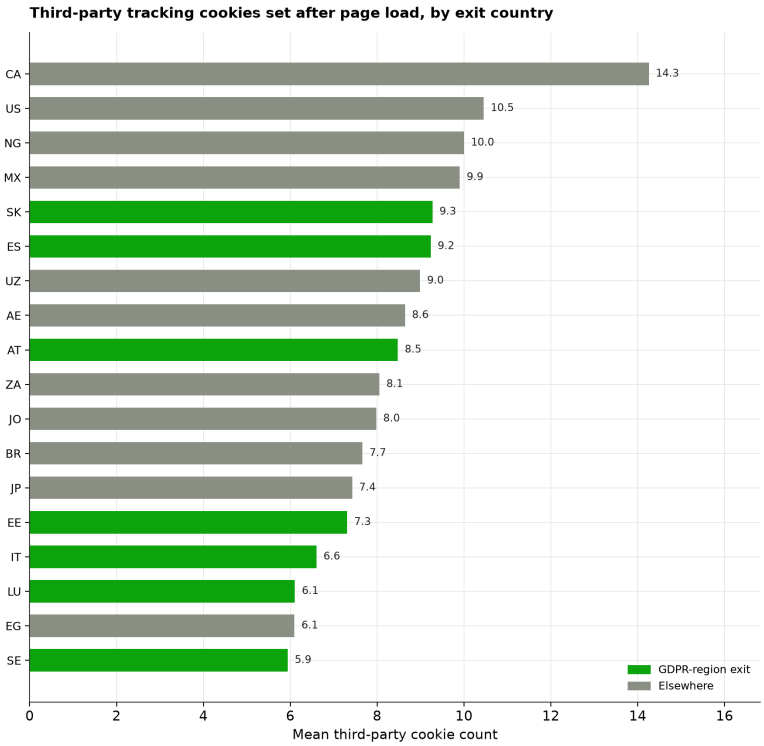

(b) Third-party tracking cookies

Figure 4.7: Website tracking and consent behavior by VPN exit country.

4.3.5 ASN Effect Versus Location Effect

Taken together, the infrastructure-isolation experiments provide limited support for the hypothesis that ASN alone determines VPN blocking. When comparing block rates across exits that share the same ASN but differ in exit country against exits that share the same country but differ in ASN, the two sweeps produced a broadly similar degree of within-group variation: neither holding ASN fixed nor holding location fixed collapsed block-rate variation to near zero, and neither produced a dramatically larger spread than the other. This suggests that, at the level of blocking outcomes specifically, ASN and location contribute comparably modest and overlapping amounts of variation, rather than one factor clearly dominating the other.

A different picture emerges when the same comparison is applied to network-path behavior rather than blocking itself. Examining which CDN edge a request was routed to reveals a much sharper asymmetry. Holding ASN fixed while varying exit country produced substantially more spread in observed CDN edge than holding country fixed while varying ASN. In other words, geographic location appears to be the dominant driver of where a request is routed within a CDN's network, while ASN plays a comparatively minor role in that particular outcome — the reverse of what a purely ASN-reputation-driven model of VPN differentiation would predict.

This divergence is informative. If ASN reputation were the main cause of VPN-related blocking, it should consistently explain differences across measurements. Instead, block rates show no clear effect from either ASN or location, while CDN routing is strongly influenced by location. This suggests that different aspects of VPN behavior are affected by different parts of the underlying infrastructure.

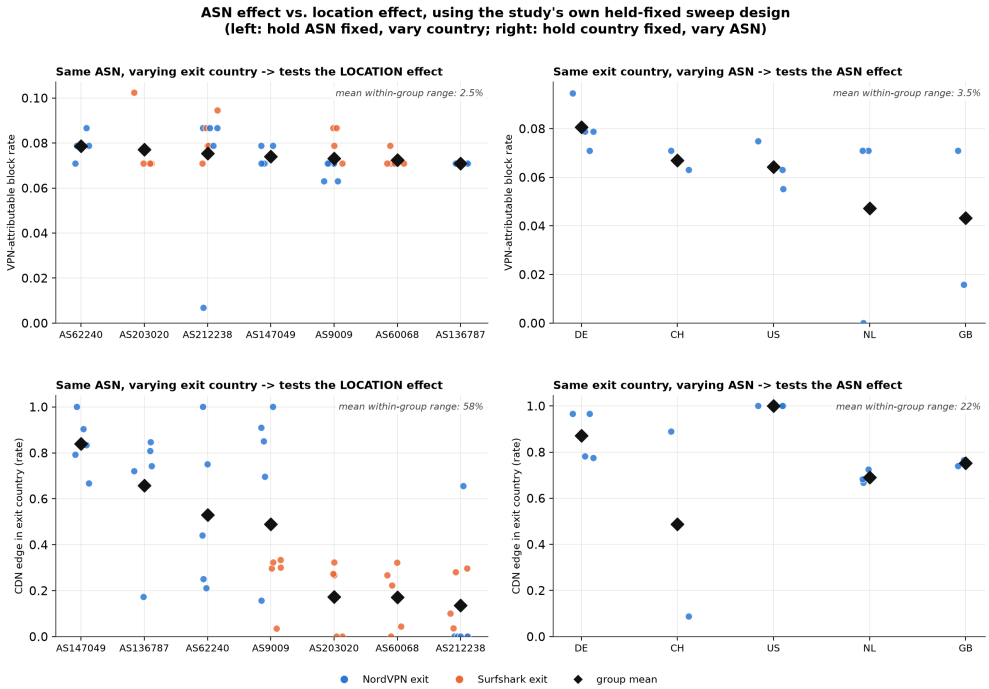

Figure 4.8: Comparison of observed variation associated with ASN and exit location across blocking and content-delivery routing measurements.

Overall, the infrastructure-isolation experiments refine the full-scale results. Different VPN exits can produce different blocking outcomes, but the larger set of exits does not show a consistent ASN- or location-based blocking effect. In contrast, location has a strong effect on CDN routing. These results suggest that VPN-related website behavior depends on several factors, including geography, ASN, routing, and site-specific policies, rather than one factor alone.

4.4 Summary of Results

NordVPN showed 1.59% VPN-attributable blocking and Surfshark 0.42%, measured across the Tranco top 10,000 domains using a single exit per provider (Appendix 7.1). Under infrastructure isolation across five ASNs and five coun-

tries (Tranco top 200), block-rate variation was similar whether ASN or location was held fixed (2.5% vs. 4.0% mean within-group range), while CDN routing varied more when location was held fixed (76%) than when ASN was held fixed (58%). Per-provider significance testing further found that most individual ASNs did not differ significantly in blocking from the other exits of the same provider (Appendix 7.3).

Chapter 5

Discussion

In this chapter, the findings are considered together to answer the central research question: to what extent can VPN blocking be explained by ASN, geographic location, or VPN provider identity? The results show that blocking varies across VPN exits, but this variation does not follow a single, consistent pattern.

The evidence instead points to a more complicated process. Websites may consider the individual exit IP, ASN, location, hosting provider, and other signals when deciding how to handle a request. The hypothesis is therefore partially supported: ASN appears to be relevant to VPN-related web behavior, but the evidence does not support it as the main explanation for VPN blocking. The following sections interpret these findings, relate them to prior work, and discuss their implications for VPN providers, website operators, and users.

5.1 Interpretation of Findings

The results show that VPN blocking is affected by the particular exit being used, but the reason cannot be isolated to ASN, location, or provider identity alone.

ASN was associated with some differences in network behavior, including CDN routing. Exits within the same ASN often showed similar CDN behavior.

Location was also examined as a possible factor in blocking. In this thesis, geographic effects were evaluated through the user experience associated with different exit countries, including differences in blocking, language or regional content, cookie-consent mechanisms, and third-party tracking behavior. These results show that exit location can influence how a website responds to a request.

5.2 Implications for the ASN-Reputation Hypothesis

The results provide limited support for the hypothesis that VPN blocking is significantly affected by ASN reputation. Block rates differed across ASNs. Even ASNs with the largest number of sampled exits and a relatively high block rate did not show a clear blocking pattern that isolated it from other ASNs.

This does not mean that ASN is irrelevant. ASN can provide websites with information about the network behind an IP address, and the experiments showed that ASN was related to other network behaviors such as CDN routing and connection success.

5.3 Implications for VPN Providers, CDNs, and Users

The findings have different practical implications for VPN providers, website operators, and users.

For VPN providers. The choice of exit infrastructure can directly affect whether users can access websites. The exit inventory also showed that NordVPN and Surfshark use infrastructure from overlapping hosting networks. This means that providers should consider the reputation and accessibility of individual exits when selecting and managing their infrastructure, rather than treating all exits as equivalent.

For CDNs and website operators. The results show that IP and ASN-based blocking can affect legitimate VPN users who share the same infrastructure with other traffic. VPN use alone does not indicate malicious activity, so broad network-level blocking can result in unintended access restrictions. Where possible, operators could use more targeted approaches that do not rely solely on the reputation of a shared IP or ASN. Privacy Pass [7] is one example of an approach that can establish trust without relying entirely on IP reputation.

For users. Not all VPN exits provide the same level of access. If a website is blocked, the problem may be specific to the exit server rather than the VPN provider as a whole. Trying a different exit may therefore restore access without requiring the user to stop using a VPN.

5.4 Limitations

Several methodological constraints affect how the results in Chapter 4 should be interpreted.

Provider scope. The study compares only NordVPN and Surfshark. These providers also share some underlying hosting infrastructure, limiting the diversity of the comparison. The findings therefore cannot be generalized to the VPN industry as a whole.

Different measurement precision. NordVPN allows direct server selection, while Surfshark required manual OpenVPN configurations. This makes Surfshark's exit selection less precise, and no complete by-country sweep was possible for Surfshark. The two providers' isolation experiments are therefore not perfectly comparable.

Small isolation sample. The isolation experiments used five ASNs, five countries, and five exits per group. This is enough to identify broad patterns but too small for strong statistical conclusions about individual groups.

Smaller domain set. The full crawl used the Tranco top 10,000 domains, while the isolation experiments used only the top 200 because of their higher operational cost. This reduces the number of blocking events and limits domain and category-level analysis.

Limited block classification. The crawler records status codes, timing, redirects, and screenshots, but not full response content or headers. It can therefore identify blocking without always determining its specific cause, such as a CAPTCHA, WAF rejection, or bot detection.

Single control vantage point. The non-VPN control used one fixed network location. Some observed differences could therefore reflect properties of the control network rather than the VPN exit alone.

Short measurement window. Data was collected over a limited period, so the results represent a snapshot of blocking behavior rather than a stable long-term pattern.

5.5 Future Work

The findings and constraints above suggest several directions for extending this work.

Greater provider and infrastructure diversity. Future work should include more diverse VPN architectures, such as owned or rented datacenter IPs, residential IPs, multi-hop routing, and decentralized relay networks. Since different providers can share infrastructure, providers should be selected based on differences in IP type, ASN ownership, and routing rather than popularity alone. Residential IPs are especially useful as a comparison because they are generally harder to distinguish from normal consumer traffic [22, 5].

Consistent exit selection. Future studies should use precise exit selection across all providers, including a complete country sweep for Surfshark. This would make ASN and location comparisons more directly comparable.

Larger samples. Testing more ASNs, countries, exits, and domains would improve statistical power and allow stronger conclusions about which infrastructure factors affect blocking.

Longitudinal monitoring. Repeating the measurements over a longer period would show whether blocking patterns remain stable or change over time. Platforms such as Censored Planet demonstrate the value of continuous internet measurement [31]; a similar approach focused on VPN exits could reveal patterns missed by a single measurement window.

External reputation data. Comparing blocked exits with external reputation services such as Spamhaus, AbuseIPDB, GreyNoise, and IPQualityScore could help test whether blocked exits are also independently flagged as suspicious.

Identifying block mechanisms. Capturing full response content and headers

would allow blocks to be classified more precisely, such as CAPTCHA challenges, WAF rules, or bot detection. This would help distinguish VPN-specific blocking from broader anti-automation measures.

Non-IP-based trust mechanisms. Future work could also examine alternatives to IP-based reputation, such as Privacy Pass [7], to explore whether legitimate VPN users can be distinguished from abusive traffic without relying solely on IP reputation.

Chapter 6

Conclusion

This thesis set out to answer a simple question: when a website blocks a VPN user, is it blocking the VPN, or the network the VPN happens to be routed through? By comparing NordVPN and Surfshark across a full-scale crawl of the Tranco top 10,000 domains, and by isolating exit servers by ASN and by location across a smaller domain set, this work found that neither factor can be isolated as the sole cause. Further, ASN ownership itself does not reliably predict blocking direction: the same hosting organization was associated with significantly higher blocking on one of its ASNs and significantly lower blocking on another.

The practical upshot is that IP and ASN-based blocking remains a blunt instrument, and that VPN "brand"- and even hosting-company ownership- is a poor proxy for the infrastructure actually driving a user's experience. Closing that gap will take broader provider comparisons, longer observation windows, and trust mechanisms that do not lean on IP reputation at all.

Chapter 7

Appendices

7.1 VPN Attribution

A domain is considered VPN-attributably denied when it is denied under the VPN condition but not denied under the control condition.

The VPN-attributable denial rate was calculated as:

$$\text{VPN-attributable denial rate} = \frac{\text{domains denied on VPN but not control}}{\text{domains present in both crawls}}.$$

7.1.1 Pipeline B – Single Exit

Definition	NordVPN	Surfshark	Control
Raw block rate (status only)	15.0%	14.0%	13.93%
Raw denied rate (block + fail)	42.2%	40.0%	39.7%
VPN-attributable, block-only	1.59% (159/10,000)	0.42% (42/10,000)	—
VPN-attributable, block + fail	3.37% (337/10,000)	1.10% (110/10,000)	—

Table 7.1: VPN-attributable blocking in the Tranco top-10,000 crawl using a single exit per provider.

7.1.2 Pipeline C – Multiple Exits

65 exits total across 13 distinct ASNs and 39 distinct exit countries. "VPN-attrib. block" excludes domains that also returned 403, 429, or 451 on the no-VPN control, as well as 61 near-universally unreachable infrastructure domains.

Provider	Exits	Connect ok	ASN mis-match	Raw block	VPN-attrib. block	Mean load
NordVPN	45	86.7%	20.5%	13.5%	1.72%	9.9 s
Surfshark	20	100.0%	42.1%	14.9%	1.96%	8.5 s

Table 7.2: Provider-level summary of the multi-exit infrastructure-isolation sweep.

7.1.3 Provider-Level Results

Splitting by provider removes cross-provider pooling: AS212238 (Datacamp) is directionally elevated for both providers but reaches significance only under Surfshark, where it accounts for a larger share of a smaller exit pool (5 of 20); under NordVPN it is diluted against a much larger exit pool. AS60068 (also Datacamp-owned, but

a different netblock) is significantly *lower*-blocking for Surfshark, underscoring that ownership/org name alone does not predict blocking direction.

Provider	ASN (org)	Exits	Mean block	Rest (same provider)	p-value	Sig.
NordVPN	AS212238 Datacamp Limited	6	2.55%	1.57%	0.099	No
	AS62240 Clouvider	3	2.22%	1.67%	0.430	No
	AS147049 PacketHub S.A.	6	2.08%	1.65%	0.097	No
	AS207137 PacketHub S.A.	1	1.67%	1.72%	—	No (n=1)
	AS136787 PacketHub S.A.	12	1.56%	1.79%	0.404	No
	AS25369 Hydra Communications Ltd	2	1.25%	1.74%	0.434	No
	AS262287 Latitude.sh LTDA	1	1.25%	1.73%	—	No (n=1)
	AS9009 M247 Europe SRL	7	1.19%	1.83%	0.012	**Yes (lower)**
	AS141039 PacketHub S.A.	1	0.00%	1.76%	—	No (n=1)
Surfshark	AS212238 Datacamp Limited	5	3.17%	1.56%	0.003	**Yes (higher)**
	AS62240 Clouvider	1	1.67%	1.97%	—	No (n=1)
	AS209854 Cyberzone S.A.	1	1.67%	1.97%	—	No (n=1)
	AS396356 Latitude.sh	1	1.67%	1.97%	—	No (n=1)
	AS203020 HostRoyale Technologies Pvt Ltd	5	1.50%	2.11%	0.127	No
	AS9009 M247 Europe SRL	2	1.25%	2.04%	0.252	No
	AS60068 Datacamp Limited	4	0.83%	2.24%	<0.001	**Yes (lower)**

Table 7.3: VPN-attributable block rate for each ASN versus the rest of that same provider's exits (Welch's t-test, one observation per exit).

7.2 ASN-Level Blocking Results

NordVPN			**Surfshark**		
Rank	ASN and network	Exit IPs	Rank	ASN and network	Exit IPs
1	AS212238 Datacamp	3,310	1	AS209854 Cyberzone	125
2	AS9009 M247	770	2	AS212238 Datacamp	53
3	AS136787 PacketHub	754	3	AS9009 M247	33
4	AS62240 Clouvider	750	4	AS9009 M247	33
5	AS147049 PacketHub	668	5	AS60068 Datacamp (CDN77)	13

Table 7.4: VPN-attributable block rates by observed exit network, including sample size and comparison with the remaining exits.

Bibliography

[1] NordVPN server list API. `https://api.nordvpn.com/v1/servers`, 2026. Accessed 2026-08-26.

[2] Not all roads lead to Rome: How VPN selection alters what we measure and infer about web infrastructure. `https://arxiv.org/abs/2605.30692`, 2026. Preprint; verify authorship before citing.

[3] Francine Biazin do Nascimento and Stefanie Roos. Measuring the accessibility of popular websites while using Mullvad VPN. Technical report, Delft University of Technology, 2021.

[4] Anupam Chander and Molly Land. United nations general assembly resolution on the right to privacy in the digital age. *International Legal Materials*, 53:727, 2014.

[5] Elisa Chiapponi, Marc Dacier, Olivier Todisco, Massimo Fangar, and Vincent Rigal. Inside residential IP proxies: Lessons learned from large measurement campaigns. *Preprint (ResearchGate)*, 2023.

[6] Cisco Systems. What is a VPN? – virtual private network. `https://www.cisco.com/c/en/us/products/security/vpn-endpoint-security-clients/what-is-vpn.html`.

[7] Alex Davidson, Ian Goldberg, Nick Sullivan, Ganesh Tankersley, and Filippo Valsorda. Privacy pass: Bypassing internet challenges anonymously. *Proceedings on Privacy Enhancing Technologies (PoPETs)*, 2018(3):164–180, 2018.

[8] William J. Drake, Vinton G. Cerf, and Wolfgang Kleinwächter. Internet fragmentation: An overview. Technical report, World Economic Forum, January 2016.

[9] Fortune Business Insights. Virtual private network market size, industry share, forecast to 2034. `https://www.fortunebusinessinsights.com/industry-reports/virtual-private-network-market-100637`, July 2026.

[10] Aniko Hannak, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson. Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the Internet Measurement Conference*, pages 305–318, 2014.

[11] Muhammad Ikram, Narseo Vallina-Rodriguez, Suranga Seneviratne, Mohamed Ali Kaafar, and Vern Paxson. An analysis of the privacy and security risks of Android VPN permission-enabled apps. In *Proceedings of the 2016 ACM Internet Measurement Conference*, pages 349–364, 2016.

[12] Eric Jardine. Tor, what is it good for? political repression and the use of online anonymity-granting technologies. *New Media & Society*, 20(2):435–452, 2018.

[13] Ruogu Kang, Stephanie Brown, and Sara Kiesler. Why do people seek anonymity on the internet? informing policy and design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2657–2666, 2013.

[14] Etienne Khan, Anna Sperotto, Jeroen van der Ham, and Roland van Rijswijk-Deij. Stranger VPNs: Investigating the geo-unblocking capabilities of commercial VPN providers. In *International Conference on Passive and Active Network Measurement (PAM)*, pages 46–68, 2023.

[15] Mohammad Taha Khan, Joe DeBlasio, Geoffrey M. Voelker, Alex C. Snoeren, Chris Kanich, and Narseo Vallina-Rodriguez. An empirical analysis of the commercial VPN ecosystem. In *Proceedings of the Internet Measurement Conference*, pages 443–456, 2018.

[16] Sheharbano Khattak, David Fifield, Sadia Afroz, Mobin Javed, Srikanth Sundaresan, Vern Paxson, Steven J. Murdoch, and Damon McCoy. Do you see what I see? differential treatment of anonymous users. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2016.

[17] Rashna Kumar, Sana Asif, et al. VPNalyzer: Systematic investigation of the VPN ecosystem. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2022.

[18] Hadi Larijani and Anna Philip. Empirical performance analysis of WireGuard vs. OpenVPN in cloud and virtualised environments under simulated network conditions. *Computers*, 14(8):326, 2025.

[19] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. Tranco: A research-oriented top sites ranking hardened against manipulation. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2019.

[20] Reethika Master and Christina Garman. Decline in IP-level blocking, 2023. As cited in: A Survey of Internet Censorship and its Measurement: Methodology, Trends, and Challenges, arXiv:2502.14945.

[21] Allison McDonald, Matthew Bernhard, Luke Valenta, Benjamin VanderSloot, Will Scott, Nick Sullivan, J. Alex Halderman, and Roya Ensafi. 403 forbidden: A global view of CDN geoblocking. In *Proceedings of the Internet Measurement Conference*, pages 218–230, 2018.

[22] Xianghang Mi, Xuan Feng, Xiaojing Liao, Baojun Liu, XiaoFeng Wang, Feng Qian, Zhou Li, Sumayah Alrwais, Limin Sun, and Ying Liu. Resident evil: Understanding residential IP proxy as a dark service. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2019.

[23] Milton L. Mueller. *Will the Internet Fragment? Sovereignty, Globalization and Cyberspace*. Polity Press, 2017.

[24] Nord Security. NordVPN: Fast and secure VPN service. `https://nordvpn.com`. Accessed: 2026-08-25.

[25] Eli Pariser. *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Press, 2011.

[26] Carlos J. C. Pena and John Evans. Performance evaluation of software virtual private networks (VPN). In *Proceedings of the 25th Annual IEEE Conference on Local Computer Networks (LCN)*, 2000.

[27] Vasile C. Perta, Marco V. Barbera, Gareth Tyson, Hamed Haddadi, and Alessandro Mei. A glance through the VPN looking glass: IPv6 leakage and

DNS hijacking in commercial VPN clients. *Proceedings on Privacy Enhancing Technologies*, 2015(1):77–91, 2015.

[28] Precedence Research. Virtual private network (VPN) market size to hit USD 534.22 billion by 2034. `https://www.precedenceresearch.com/virtual-private-network-market`, November 2025.

[29] Security.org. 2026 research into VPN usage (updated statistics). `https://www.security.org/vpn/statistics/`, March 2026.

[30] Rachee Singh, Rishab Nithyanand, Sadia Afroz, Paul Pearce, Michael Carl Tschantz, Phillipa Gill, and Vern Paxson. Characterizing the nature and dynamics of Tor exit blocking. In *Proceedings of the 26th USENIX Security Symposium*, pages 325–341, 2017.

[31] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. Censored planet: An internet-wide, longitudinal censorship observatory. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 49–66, 2020.

[32] Surfshark B.V. Surfshark: Fast, simple, affordable VPN. `https://surfshark.com`. Accessed: 2026-08-25.

[33] Michael Carl Tschantz, Sadia Afroz, Shaarif Sajid, Shoaib Asif Qazi, Mobin Javed, and Vern Paxson. A bestiary of blocking: The motivations and modes behind website unavailability. In *8th USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2018.

[34] Windscribe. 55 VPN statistics & industry trends (2023–2025). `https://windscribe.com/blog/vpn-statistics-trends/`, February 2025.