

An Empirical Analysis of the Commercial VPN Ecosystem

Mohammad Taha Khan*
UIC

Joe DeBlasio*
UC San Diego

Geoffrey M. Voelker
UC San Diego

Alex C. Snoeren
UC San Diego

Chris Kanich
UIC

Narseo Vallina-Rodriguez
IMDEA Networks Institute
ICSI

ABSTRACT

Global Internet users increasingly rely on virtual private network (VPN) services to preserve their privacy, circumvent censorship, and access geo-filtered content. Due to their own lack of technical sophistication and the opaque nature of VPN clients, however, the vast majority of users have limited means to verify a given VPN service's claims along any of these dimensions. We design an active measurement system to test various infrastructural and privacy aspects of VPN services and evaluate 62 commercial providers. Our results suggest that while commercial VPN services seem, on the whole, less likely to intercept or tamper with user traffic than other, previously studied forms of traffic proxying, many VPNs do leak user traffic—perhaps inadvertently—through a variety of means. We also find that a non-trivial fraction of VPN providers transparently proxy traffic, and many misrepresent the physical location of their vantage points: 5–30% of the vantage points, associated with 10% of the providers we study, appear to be hosted on servers located in countries other than those advertised to users.

ACM Reference Format:

Mohammad Taha Khan*, Joe DeBlasio*, Geoffrey M. Voelker, Alex C. Snoeren, Chris Kanich, and Narseo Vallina-Rodriguez. 2018. An Empirical Analysis of the Commercial VPN Ecosystem. In *2018 Internet Measurement Conference (IMC '18)*, October 31–November 2, 2018, Boston, MA, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3278532.3278570>

1 INTRODUCTION

Numerous large-scale data breaches [18] and increasingly prevalent reports of censorship [33] have Internet users seeking out tools to help preserve their privacy online and circumvent potential censorship. Virtual private network (VPN) services in particular play an increasingly integral role in securing Internet freedom. Likely spurred by the large number of media articles recommending their use [11, 12], VPN usage has grown dramatically in recent years. According to a recent market research report [32], commercial VPNs are currently a 15-billion dollar industry expected to grow by 20% by 2022.

*Authors Khan and DeBlasio contributed equally to this work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '18, October 31–November 2, 2018, Boston, MA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5619-0/18/10...\$15.00

<https://doi.org/10.1145/3278532.3278570>

Originally developed as a technology to privately send and receive data across public networks, VPNs are now marketed broadly as a privacy-preserving technology that allows Internet users to obscure not only their traffic, but also their personal information, such as their web browsing history, from third parties including Internet service providers (ISPs) and governments alike. VPNs are widely used by individuals not only to preserve their privacy, but also to evade Internet censorship—whether implemented by governmental organizations or by individual content providers who geo-filter content to, *e.g.*, enforce copyright or licensing agreements. Compared to sophisticated tools like Tor [34], commercial VPN services purport to provide individuals a turnkey solution to achieve their privacy goals and to access blocked content [17, 23, 35]. As a result, many users often prefer VPNs over free-but-complicated services like Tor due to performance claims and perceptions of enhanced usability.

Unfortunately, while many services claim to operate robust and secure infrastructure and ensure user privacy by not logging data, the reality is that the VPN ecosystem is highly opaque. This state of affairs is further aggravated by the lack of practical tools or independent research that systematically audits these security and privacy claims, especially with respect to information leakage and traffic manipulation. Anecdotally, some VPN providers are known to sell customer data to third-party data brokers or manipulate customer traffic. Previous studies in the mobile space [13] reveal some VPN providers acting maliciously or presenting worrisome vulnerabilities. These findings were followed by FTC investigations into Hotspot Shield allegedly intercepting and manipulating user traffic [4].

The absence of independent and peer-reviewed evaluation of VPN services leaves privacy- and security-conscious users little choice but to resort to blogs, word of mouth, and review websites when shopping for VPN services. Unfortunately, we find that many of these websites are supported by affiliate programs, suggesting they are unlikely to provide unbiased opinions. As a first step toward a rigorous, third-party evaluation of VPN service claims, we develop a generic test suite and apply it to a diverse set of 62 VPN providers. We seek not only to highlight the issues surrounding the lack of transparency in the commercial VPN ecosystem, but also develop an active-measurement methodology that others can use to evaluate and audit arbitrary VPN services.

In addition to widespread issues with transparency, marketing, and systematic traffic leakage due to poor security defaults, we find evidence that approximately 10% of VPNs are intercepting and/or manipulating user traffic. While (visible) interception is less frequent than reports from other, similarly-situated industries [13], it is important to note that VPN operators can more easily passively monitor traffic in ways that are difficult to detect from the standpoint

of an end user. Hence, our findings regarding traffic tampering and monitoring are at best a lower bound on their actual prevalence.

In addition to factors that compromise user privacy, we also find that many VPNs fail to deliver on their promises of geographic diversity. VPN providers frequently advertise that their service can make user traffic appear to originate from a selection of distinct vantage points, typically spread across different countries. Our study shows that the practice of ‘virtualizing’ vantage points—*i.e.*, physically placing VPN end points in one (presumably more accessible) country and then working to trick IP geo-location services into believing that the vantage point is in another country—is far more prevalent than VPN services let on. We find that 10% of the providers in our study appear to misrepresent the location of one or more of their vantage points, with 5–30% of all vantage points (depending upon the source of ground truth) located in a different country than advertised. In the most extreme case, we find a VPN provider claiming vantage points in more than 190 countries, yet hosting servers in what appears to be fewer than 10 distinct data centers.

2 BACKGROUND

Before describing our measurement methodology, we begin with a brief overview of the risks facing VPN users, and elaborate upon the need for independent verification.

2.1 Risks of exposure

Users’ fears of surveillance are well founded, not just internationally, but even within the United States. In March of 2017, the US Congress enacted a Congressional Review Act (CRA) to repeal privacy rules developed by the Federal Communications Commission (FCC) [11, 10]. The rules would have required ISPs to seek permission from customers for collecting and sharing sensitive personal information such as their Internet browsing history. In response to the repeal of the rules, public concern has prompted privacy advocates and other responsible entities to point to VPNs as a viable way to regain some control over their private information.

Unsurprisingly, as of September 2018, many countries regulate VPN use, and some, including China, Iran, Russia, Syria and Egypt, attempt to block VPN services in a variety of different ways [3]. For instance, Russia, Iran and China regulate VPN usage by only allowing use of government-approved providers [12]. Other countries, like the United Arab Emirates, only allow the use of VPN services by certain approved institutions like banks [3]. Hence, any vulnerabilities present in commercial VPN services—intentional or otherwise—that expose its users to surveillance may subject them to potential fines, criminal prosecution, or worse.

VPN services are also widely used to access geo-filtered content and download copyrighted material [15]. For instance, Germany has very strict laws regarding illegal sharing of copyrighted content, and its supreme court has empowered to require ISPs to hand over contact information of Internet users who infringe the law by uploading or downloading copyrighted content without authorization. As a result, many German users employ VPN services to obscure their IP address when using peer-to-peer systems. On the flip side, popular content providers like Hulu and Netflix actively detect and block users from outside their authorized service areas from connecting to

their services. In turn, VPN providers also implement features that evade server-side content blocking [40].

2.2 Questionable quality standards

While users flock to commercial VPN services, the effectiveness of most tools remains unclear: Only a handful of VPN providers offer users an audit trail or the ability to verify that their traffic is not being leaked. Despite VPNs having been around for more than a decade, Tunnelbear was the first VPN provider to release a third-party security audit in 2017 [37]. In addition, there is also a dearth of user-friendly transparency tools which allow users to independently evaluate the claims made by individual VPN services. NordVPN has a DNS-leakage detection test, for instance, but it has limited scope [25]. Not only is there a lack of widespread positive evidence of quality, there is in fact anecdotal evidence to the contrary. Recently, the Center for Democracy and Technology (CDT) filed a complaint against HotSpot Shield VPN [4], an Android VPN service offered by AnchorFree Inc. that was actively redirecting user traffic to partner websites. This incident further supports the need for a comprehensive and transparent evaluation of commercial VPNs.

Moreover, the fierce competition between providers, combined with the lack of objective metrics to evaluate them and an unregulated Internet market, lead to considerable deception to attract clients. Many Internet users rely on online ratings and ranking sites when selecting and shopping products, and VPN services are no exception. It is troubling, then, that several top-ranked VPN review websites are currently supported by VPN affiliate marketing and services. With strong economic incentives to funnel users to particular VPNs, it is hard to be confident in the impartiality of the review websites. For instance, VPNmentor [41], a popular VPN review site, lists over 250 VPN services and claims impartiality: “Our reviews are written by users themselves, and are not influenced by VPN companies.” However, none of the listed VPNs have a review score below 4 out of 5. As a result, users blindly trusting the information provided by rating websites are likely to select a VPN service that provides a handsome affiliate payoff, but fails to meet the security and privacy expectations of the user.

3 DATA SOURCES

While there are numerous VPN services available for access on the Internet, it can be challenging to evaluate each and every single one of them. Our goal in this study is to perform a broad analysis of VPN services which truly reflects the state of the ecosystem, without having to exhaustively evaluate every one. Hence, we have tried to select services that have a larger impact on the security and privacy implications of individuals. We created a comprehensive list of VPN services by extracting VPN names from review websites, as well as through personal outreach on email lists and forums, and then selected a subset to study based on explicit review criteria.

We collected our candidate set of services as follows:

- **Popular review sites:** Popularity was one of our primary metrics to ensure sound and reasonable selection of VPN services. We extracted the URLs of the first 50 Google search results when using the search keyword “*top VPN services*”. The results included links to VPN websites as well as review-based web services with listings of top VPN providers. We

Website	Affiliate Based Link
360topreviews.com	✓
bbestvpn.com	✓
best.offers.com	✓
bestvpn4u.com	✓
freedomhacker.net	✓
ign.com	✓
pcmag.com	✓
pcworld.com	✓
reddit.com	✗
securethoughts.com	✓
techsupportalert.com	✓
thatoneprivacysite.net	✗
tomsguide.com	✓
top10fastvpns.com	✓
torrentfreak.com	✓
trustedreviews.com	✓
vpnfan.com	✓
vpnmentor.com	✓
vpnsrus.com	✓
vpnservice.reviews	✓

Table 1: Websites used to populate categories for the aggregated VPN list along with their affiliate marketing status.

extracted the VPN names in the search results and performed additional crawls on the review sites. The intuition behind this approach was to imitate users performing an online search for a reliable VPN option. Unsurprisingly, almost all of the top sites participate in some form of affiliate marketing. So the rankings here must be taken with a grain of salt, but they offer the best public metric for popularity available.

- **Reddit crawl:** The subreddit channel for VPNs [30] is an active forum where individuals seek out recommendations and discuss general issues related to VPN setup and configuration. In July 2017, we performed a complete crawl of the two main ‘mega-threads’ about VPN recommendations. In addition we performed a secondary crawl to extract VPN names from regular posts on the VPN subreddit within the past three months.
- **Personal recommendations:** To ensure we included VPNs common in censoring countries, we reached out on the Open Technology Fund [27] mailing list. We specifically requested feedback from individuals from countries implementing censorship who recommended privacy enhancing mechanisms to local users and were aware of the VPN services that were regionally popular.

Table 1 shows the complete set of review websites we considered, most of which are affiliate-based. Two websites, VPNMentor [41] and The One Privacy Site [26] also have basic descriptive details of VPN services such as their costs, operational features, number of vantage points and even user reviews in various languages. This additional information allowed us to refine our selection process into independent categories and diversify our selection subset. In particular, we focused on the following features:

VPN Selection Category	# of VPNs
Popular Services (from review websites)	74
Reddit Crawl	31
Personal Recommendations	13
“The One Privacy Site”	
Cheap & Free VPNs	78
“VPN Mentor”	
Multiple Language Reviews	53
Large Number of Vantage Points	58
Others	45
Total Selected	200

Table 2: Number of VPNs extracted from each source of providers. Note that sources overlap substantially.

- **Cheap and free services:** VPN subscription price is a major factor individuals take into account when selecting a VPN service. Using the crawled details we shortlisted VPNs which had a monthly subscription cost of less than \$3.99. This approximately equated to the lower quartile of the price distribution of the collected set. While we focused on the collection of commercial VPN services, we also augmented our list with VPNs that specifically had a free-to-use version which was not limited by time-based trial and could be used indefinitely with restricted features.
- **Large number of vantage points:** One of our selection criteria was for services that claim to have a large number of vantage points. Using our seed lists, we selected the ones which provided connectivity via 30 or more countries. This category was also formed on the basis of user intuition that individuals are likely to select services with more geographic spread.
- **Multiple language reviews:** VPNMentor [41] was the only website which contained user-based reviews for VPN services. This feature enabled us to extract VPNs which had reviews written in multiple languages. We created this category in particular to ensure global diversity in our collection process.
- **Others:** This category consisted of VPNs which did not have a significant impact on our motivation, yet were still noteworthy to take into consideration. These included services which provided connectivity over multiple tunneling protocols such as IPSec, L2TP and SSH. We also shortlisted VPNs which provided multiple simultaneous connections.

Table 2 lists the number of VPNs that we collected, as well as the number that fell into each (non-exclusive) category. As these individual lists had overlap, we took their union to create a comprehensive merged list with 200 unique commercial VPN services.

4 ECOSYSTEM ANALYSIS

Next, we methodically mined information from the websites of the VPN services in the list. We performed this process in a partially automated way, followed by substantial manual effort for post-processing and inspection to verify correctness. In some instances,




Figure 1: Geographic distribution of VPN business locations

when a certain feature was not mentioned on the website, we also performed active outreach on the VPN website contact forms. Based on the data collected from these VPN services, we now provide a systematic overview of the ecosystem along with empirical details of the features collected, and highlight notable trends among them.

The Emergence of VPN Services: From the initial list of VPN services, we selected the top 50 VPNs from the popular services category. 15 of these VPN services were founded between the years 2005 and 2013. Some of the oldest providers were HideMyAss¹, IPVanish, StrongVPN and Ironsocket, all of which were founded in 2005. Overall, 45 (90%) of the VPN services were founded after 2005. These founding dates suggest that the growth of VPNs was mainly motivated by the increased number of Internet users seeking anonymity and circumvention tools during that era.

VPN Business Locations: We also collected country-level details for the VPNs using their business locations listed on their websites, or by reaching out on their contact forms. Figure 1 shows the country-level geographic distribution where VPNs claim to be based. A majority of the VPN services are located in countries with no censorship laws. These include the US, UK, Germany, Sweden and Canada. Only two VPNs, FreeVPN Ninja and Seed4.me, claimed to be in China. Recently, FreeVPN Ninja was discontinued, and it was mentioned on the Seed4.me blog that it is now blocked due to the lack of compliance with the Chinese government VPN law [31]. Interestingly, we also observed a handful of services having locations in Seychelles and Belize, both of which are relatively small countries that have experienced Internet regulation by their governments. Another noteworthy case is NordVPN, which has 1,665 servers alone in the US while the company itself is based in Panama. They primarily claim to operate in this way to evade secret warrants, such as government subpoenas or data acquisition letters. For added transparency, NordVPN also issues a “warrant canary”, a public statement informing users that the company has not received any legal process requests.

Marketing and Affiliate Strategies: VPN services use multiple strategies to market themselves. They generally brand themselves as security, anonymity and Internet freedom tools. To attract users,

¹Later acquired by Avast Software Inc. in 2015.




Figure 2: Claimed server counts of VPN services

they claim various features on their websites such as the number of servers, their supported tunneling protocols, and their encryption standards. One common term seen across various VPN websites was “military grade encryption”, a marketing term used for AES-256 bit encryption. Hotspot Shield was one of the popular VPN services to claim this. Our data collection highlighted that VPN services also heavily make use of popular social media for marketing. 126 VPNs had a Facebook page and 131 had a Twitter account. For outreach through review websites, VPNs make use of affiliates. In our list, 88 of the VPNs had an affiliate program. Preliminary investigation revealed that VPN services had separate login portals for affiliates and the commissions varied based on the credibility and popularity of their affiliate partners. This aspect specifically highlights the opacity of the operation and marketing of VPN providers.

Transparency Practices of VPNs: To obtain a baseline for the transparency of these VPN services, we looked for how many had a privacy and an acceptable usage policy on their websites. 25% (50) of the VPN services in the list did not have a link to their privacy policy. 42% (85) of the VPN providers also did not provide a terms of service. As a quick measure, we examined the lengths of the privacy policies and found wide variation, ranging from a privacy policy as small as 70 words to a maximum of 10,965 words, with an average of 1,340 words. We also explored whether VPN services mentioned anything specific about their logging policies on their homepage or in their privacy policies. This information was heterogeneous and was either embedded in the privacy policies, or as a branding feature on the VPN homepage. Our analysis found only 45 VPN services explicitly claiming a “no-logs” policy. Overall, this analysis suggests that VPN providers are greatly heterogeneous in terms of transparency. While the popular services make genuine transparency efforts, due to the lack of regulation in the VPN industry, the VPNs in the tail of the popularity distribution show contrasting trends and unprofessionalism. Instances of such attitudes were evident in the irregularity of their privacy policies and the absence of terms of service for a significant number of VPNs.

VPN Vantage Points Infrastructure: One of the key factors users take into account when selecting a VPN service is the diversity of its vantage locations or servers. VPN services mention the number of vantage points on their websites as a branding feature. Figure 2 shows the distribution of claimed number VPN vantage points. 80% of the VPNs have 750 vantage points or less. The more popular VPNs such as NordVPN, Private Internet Access, Hotspot Shield and TorGuard mentioned having vantage points in the 2000 to 4000 range.

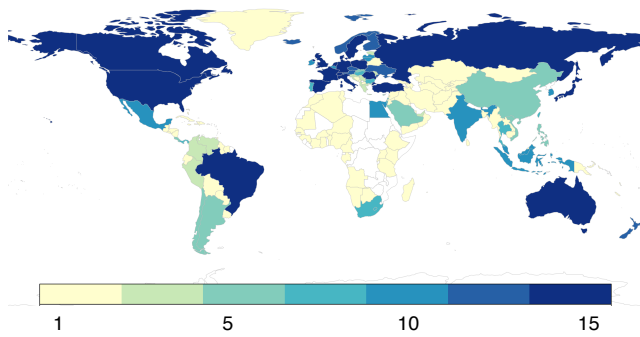


Figure 3: Geographic distribution of vantage points for the top 15 popular VPN services




Figure 4: Popular accepted payment methods

It is important to note that, while VPN services claim large numbers, their corresponding clients only provide connection options on a country-based granularity, which makes the verification of actual servers non-trivial. Furthermore, Figure 3 shows the geographic heat map of countries where the popular VPN providers have servers. Countries without censorship in North America and Europe are likely locations for more VPNs to have vantage points. Notably, HideMyAss VPN claimed vantage points in the Internet-censored regions of the Middle East, such as Iran, Saudi Arabia, and even North Korea. We validate these claims through our geo-location based tests in Section 6.4.

Subscription Models: Commercial VPN services offer users various subscription models. Table 3 shows the average, minimum and maximum rates for the subscription types offered by the VPNs. After the monthly baseline, the second most popular subscription plan offered was the annual subscription, which approximately costs half the monthly rate. 19 of the 200 VPN services had subscription offers for more than a year. These durations included two-year, five-year, and even lifetime subscriptions. For instance, CrypticVPN and HideMyIP offered lifetime access for only \$25 and \$35, respectively. We also inspected if VPN services had refund policies and trial versions. 45% of the VPNs had a free or a trial version, which was limited in time, usable bandwidth, or the available server locations. Refund policies ranged from 24 hours up to 60 days, while the seven-day full refund was the most commonly offered policy (40% of the services).

Subscription	# of VPNs	Monthly Cost (\$)		
		Min	Avg	Max
Monthly	161	0.99	10.10	29.95
Quarterly	55	2.20	6.71	18.33
6 Months	57	2.00	6.81	16.33
Annual	134	0.38	4.8	12.83

Table 3: Monthly subscription costs across various VPN subscription models




Figure 5: Tunneling technologies used by VPN services

Payment Methods: We explored the various acceptable payment methods offered by commercial VPN providers. Overall, 61% of the VPN providers accepted credit cards, 59% of them had online payment options such as Paypal, and 46% of them also accepted cryptocurrencies. 32% of VPN providers did not accept credit cards but took both online payments and cryptocurrencies.

Figure 4 shows the top three payment options from each category. In general, VPN services accept a variety of payment options, allowing users from different countries to have viable methods to pay. Among cryptocurrencies, Bitcoin was by far the most popular. Services accepting cryptocurrencies distinctly marketed themselves as accepting anonymous payment methods to attract more users.

Platform Support: From the selected VPN services, we looked at support for various operating systems and mobile platforms. 87% provided support for both Windows and OSX. Linux support was slightly less popular, with 61% of the VPNs supporting it. The setup mechanisms provided by VPN were either in the form of standalone clients, or configuration files with user instructions on how to configure the VPN tunnel via third-party software such as OpenVPN or Tunnelblick [28, 38]. Mobile device support was also a fairly common feature. 56% of the VPNs had released both an Android and an iOS app. Some VPN providers were only browser-based extensions; while we consider them in our ecosystem analysis, we exclude them from our active VPN testing discussed in Section 5.3.

Security Features: A primary feature of a VPN service is the underlying protocol that they use. The ability to support multiple tunneling protocols is used as a marketing feature and is mentioned on VPN

websites. Figure 5 shows the distribution of the tunneling protocols supported across multiple VPN services. The majority of VPN providers support both OpenVPN and PPTP tunneling. While these protocols are inherently secure, it is important to note that misconfigurations in VPN clients can still cause traffic leakage to occur. We discuss and evaluate leakage of VPN traffic in Section 6.5. As an added security feature, 18 of the VPNs mentioned having a “kill-switch”. This functionality prevents any traffic leakage in case of a tunnel failure. We explore the default state of kill-switches in VPNs and evaluate the robustness of VPNs in cases of tunnel failures. Interestingly, 10 of the VPN services also mentioned providing VPN over Tor as a service. This feature allows users to route their encrypted VPN traffic over the Tor network. While this feature has the benefits of preventing against malicious Tor exit attacks and reducing the risk of VPN logging, it also has considerable performance tradeoffs.

P2P-based VPN Services: One common VPN use case for regular Internet users is to download P2P torrents. Individuals connect via VPNs to prevent their ISPs from detecting any P2P traffic and to hide their IP addresses. 64 VPN providers in our list mentioned allowing torrents over their network. While these VPNs mentioned that they did not promote the illegal access of copyrighted content, it was unclear how they dealt with DMCA notices they might receive.

5 METHODOLOGY

This section describes the criteria used to select the subset of VPN services we experimentally evaluated, along with the details of our VPN testing infrastructure.

5.1 VPN Subset Selection

Testing VPN services is a time-intensive task. For each service, one must individually create and verify an account, optionally pay money for the service, install the software, and verify its correct operation. VPNs with custom software (rather than the ones which use a standard VPN protocol client) require manual intervention between each vantage point measured, as the user must interact with the software to connect/disconnect from each vantage point. In addition, our full test suite takes approximately 45 minutes to execute on each vantage point, and there is limited ability to parallelize due to the limited number of test machines and their computational infrastructure.

Because of manual labor, as well as the financial costs, it was not operationally feasible to perform measurements on all 200 VPN services we selected. To prioritize a subset of these VPNs, we defined a set of features listed below which allowed us to create a stratified sample of the ecosystem, balancing a broad selection of VPNs with a realistic time and monetary budget:

Popular VPN Services: To maximize coverage of current VPN users, we used the popularity metric described in Section 3 to select the top 15 VPN services for evaluation.

VPNs with Free or Trial Versions: Several of the VPN services we enumerated included free or trial options. We chose 30 VPN services from this category. This decision was primarily motivated by free/trial versions being popular in certain countries, while it also simplified our sign-up process.

Random Selection: We also randomly selected 16 VPNs to diversify our selection set.

Using these features, our final VPN list consisted of 62 VPN services, including the ones which were chosen arbitrarily.

5.2 Testing Methodology

To ensure accurate and correct data collection for these VPN services, we tested them using macOS virtual machines running across several systems. Testing within a VM allowed us to maintain complete isolation between VPNs by using a freshly-restored macOS instance for each.

We selected macOS as our primary platform as it provided a tradeoff between the authors’ experience in writing for UNIX/Linux systems and the prevalence of client software for major desktop operating systems. In instances where no client software was available but OpenVPN configurations were, tests were automated. To test each VPN, we first registered for an active subscription and downloaded the client. Performing a full test on a single vantage point for a VPN took approximately 45 minutes to complete inside of the VM environment. The suite logs results for each experiment as well as traffic traces for passive analysis. As testing from every single vantage point did not scale, we selected approximately five vantage points for each manually-evaluated VPN, attempting to maximize geographic diversity. Manual and automated collection combined, we collected data from 1046 vantage points.

Several VPN servers disconnected during the testing phase and hence required partial re-collection of data. A general trend we observed across VPNs was the irregularity in vantage point connectivity. While we were typically able to connect to VPN vantage points in North America and Europe, there was far lower reliability when connecting through vantage points in the Middle East, Africa and South America. Multiple VPN services we initially selected also had buggy or defunct clients that prevented us from testing at all. In these cases, we replaced them with others from our list.

5.3 Testing Infrastructure

Here we describe the tests that we implemented to evaluate the VPN services. Our tests fall into three main categories: tests that identify active manipulation or surveillance by the VPN provider, tests which explore the infrastructure of VPNs (for instance by estimating geolocation of vantage points), and finally tests that identify unintentional traffic leakage.

5.3.1 Traffic Interception and Manipulation Tests. To detect any evidence of manipulation or observation on the part of the VPN provider, our experiments test the following scenarios:

DNS Manipulation: This test makes DNS requests to several popular hosts, and then investigates whether any of the returned IPs look suspicious, as determined by comparing the results against Google Public DNS. Any differences are filtered by investigating the WHOIS records of the IPs returned by the non-Google server, looking for owner information. This test relies on a small number of hosts (whose WHOIS records conform to the format expected) and relies on a human to investigate suspicious-looking results. It also assumes that DNS manipulation will only occur via VPN-provided

DNS servers, and that the VPN will not spoof results from Google’s public DNS.

DOM and Request Collection: This test loads the homepage of 55 sites in a Selenium-controlled Chrome session, and saves logging information regarding what resources are loaded and why, along with a complete copy of the DOM and other debugging information. We specifically chose domains which do not upgrade requests to HTTPS to maximize the opportunity for manipulation. In addition, they also cover a variety of potentially sensitive categories (including politics, pornography, government websites, defense contracting, etc.).

Two domains in our list act as ‘honeysites’ akin to the methodology used by Tsirantonakis, *et al.* [36] and Ikram *et al.* [13]. These sites provide static DOM content to provide an opportunity for traffic modification by VPN providers. One site contained ad inclusion code from major ad networks to allow for easier ad injection, but using invalid publisher identifiers to ensure no impact to the ad network or advertisers.

This test accounts for the majority of the overall test suite’s execution time due to the considerable overhead of loading complete pages in Chrome. As a result, we had to limit our collection size to make the problem tractable.

TLS Interception and Downgrade Detection: This test augments the collection in the previous test by collecting data on the same 55 hosts, along with more than 150 additional hosts. This test performs two steps: first, it directly negotiates a TLS connection with the end host, validates the certificate, and stores the certificate for subsequent processing. Second, it loads the website of each hostname, starting via HTTP, and follows all HTTP redirects, recording response headers and URLs loaded. This test is well-situated to reveal malicious redirects and TLS stripping.

For both this test, and the previous one, we periodically collected ‘groundtruth’ (known-unmodified) data from a university IP several times per day during our study. This dataset formed the whitelist to identify inconsistencies.

5.3.2 Infrastructure Inference Tests.

Recursive DNS Origins: This test resolves a unique timestamped and tagged hostname under a domain name whose nameserver is configured to store records of where requests come from for further investigation.

Ping and traceroute data: This test collects pings and traceroutes to anycasted public DNS resolvers (Google Public DNS and Quad9) and well-distributed DNS roots (D, E, F, J, L). It also collects ping data for 50 RIPE Atlas anchors. We use this data in Section 6.4 to infer whether an endpoint is ‘virtually’ or physically located in a claimed location.

Geolocation via Google API: The Google Maps API allows geolocation of IPs making API requests to Google [9]. This test collects both the coordinates and reverse geo-coding from their API.

5.3.3 Leakage Based Tests. Most VPNs are built upon commonly-used and understood protocols largely considered secure, however misconfigured services can cause some traffic to be

exposed. These set of tests focus on identifying instances of such leakage.

DNS Leakage: This test makes a series of predetermined DNS requests both to the system’s default DNS server and to public resolvers. It collects traffic on the primary (non-VPN) interface and checks whether there are any DNS packets exiting outside the VPN interface.

IPv6 Leakage: As IPv6 is still early in adoption, most VPN services provide only IPv4 support. Our IPv6 leakage test looks for leakage by attempting to directly communicate with several popular websites with IPv6 addresses while capturing traffic on the non-VPN interface. Since IPv6 traffic should also be routed via the VPN, the test scans through the collected non-VPN interface traffic for IPv6 requests to the websites under test, and triggers if any are detected.

Recovery from Tunnel Failure: This test inspects how VPN services handle unexpected tunnel failure as caused by, for instance, temporary loss of connectivity to the VPN provider. Even if a VPN re-establishes its connection automatically, if during an outage traffic is temporarily routed outside of the VPN, the VPN provider’s security guarantees are undermined, and thus a VPN provider should ‘fail closed’. This test artificially induces a tunnel failure by firewalling all outbound connections from the hardware interface except to a fixed set of hosts. It then repeatedly attempts to contact these hosts over the course of several minutes. Failing behavior is being able to contact the hosts during a VPN failure.

This test catches many but not all VPNs that fail-open, and thus should be considered a conservative estimate. The test must decide how long to wait to allow a VPN to realize that its connection has failed. Our test used a three-minute blocking window.

5.3.4 VPN Metadata and packet captures. In addition to test-specific data, we also collect a great deal of general configuration, addressing and routing information, and packet captures to facilitate analysis, investigation and debugging. This information includes routing and ARP tables, interface lists, configured DNS resolvers, and pings to any /32 IPv4 routes. We collect this data to facilitate our own investigation into any anomalies detected.

Our normal testing also collects packet captures on the hardware interface. We subsequently analyze this traffic to detect non-VPN-traversing leakage, and to detect whether the VPN service is providing our IP address as an additional vantage point in their network. For simplicity, we focus on identifying unexpected DNS requests to identify P2P traffic.

6 RESULTS

In this section, we investigate the runtime and network behavior of the 62 selected VPN services (Section 5.1) using the tests described in Section 5.3. Appendix A lists the complete set of VPN services we evaluated.

6.1 Traffic Manipulation

Overall, most VPN providers showed little evidence of traffic manipulation or tampering. Most manipulation we witnessed was a result of a country-level censorship, rather than VPN-level malicious intent.

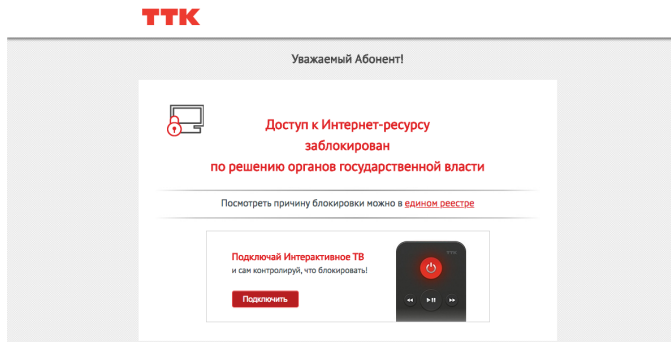


Figure 6: TTK redirection when visiting blocked content in Russia.

These results contrast with the picture shown by previous research efforts on Android VPN apps [13] and open web proxies [36], in which traffic manipulation was a common practice.

Because so many of the VPN services were highly unreliable, we were unable to detect blocking via traditional mechanisms like TCP RSTs or dropped traffic since the behavior aliases significantly with low-quality connections.

6.1.1 URL Redirection. One form of redirection our tests detect is when unauthenticated HTTP requests bound for a site encounter one or more HTTP redirects to an unrelated domain. We considered two subdomains to be related if they shared a registered domain (according to the public suffix list [21]), their registered domains differed only by public suffix (e.g., `http://a.example.com` to `http://b.example.org`), or if they were manually determined to be related. Table 4 enumerates every instance of this form of redirect that we detected. In every instance, the blocking was the result of upstream blocking based on country-level censorship in Turkey, South Korea, Russia, Netherlands and Thailand.

In each instance, we detected redirection of sensitive domains (typically via HTTP 302 redirections) only on endpoints claiming to be in their respective countries. The types of sites experiencing the highest level of blocking were pornography (Turkey, South Korea, Thailand, Russia) and file sharing (Turkey, Russia, Netherlands). Additionally, in our dataset, Turkey blocked Wikipedia and Russia blocked `jw.org` and `linkedin.com`. In the case of Russia, the redirection is enforced by the ISP following government and legal requirements as shown in Figure 6.

6.1.2 TLS Downgrades. We also investigated TLS downgrades in which requests sent to `http://foo` normally are directed to `https://foo`. We found more than a dozen instances where sites provide, for instance, HTTP 403 responses from the initial non-TLS-based page load of services across dozens of VPN providers and vantage points. These services appear to be trying to systematically block known VPN traffic. We also found a few additional instances of country-level upstream blocking of the same sites as in the previous section in which sites were similarly provided with a HTTP 403, or an HTTP 200 response code with empty or small bodies on blocked pages. We take these signals as validation of our technique, but we found no VPN providers systematically stripping TLS connections.

Destination Domain	VPNs	Country
<code>http://195.175.254.2</code>	8	Turkey
<code>http://[www.]warning.or.kr</code>	5	South Korea
<code>http://fz139.ttk.ru</code>	4	Russia
<code>http://zapret.hoztnode.net</code>	2	Russia
<code>http://warning.rt.ru</code>	1	Russia
<code>http://blocked.mts.ru</code>	1	Russia
<code>http://block.dtl.ru</code>	1	Russia
<code>http://blackhole.beeline.ru</code>	1	Russia
<code>https://www.ziggo.nl</code>	1	Netherlands
<code>http://213.46.185.10</code>	1	Netherlands
<code>http://103.77.116.101</code>	1	Thailand

Table 4: Destination domains of URL redirections.

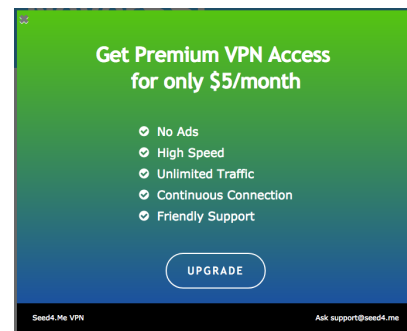


Figure 7: Advertisement for premium services injected by Seed4.me trial VPN service.

6.1.3 Traffic Injection/Modification. When comparing domains accessed by our Chrome browser instance visiting our honeysites against the expected whitelist, we identified one instance of a VPN provider systematically injecting content into pages. In this case, our test used a free trial account of Seed4.me. Pages loaded under this VPN had HTTP pages injected with custom JavaScript hosted on a subdomain of the VPN provider’s site to inject an overlaid advertisement, shown in Figure 7.

6.2 Traffic Monitoring

We used two approaches to infer whether VPNs might be monitoring user traffic.

6.2.1 Header-based Proxy Detection. Our header modification test looked for proxies that do not necessarily announce themselves by looking for differences between what data is sent in an HTTP request by our testing infrastructure and what the server sees. We detected a transparent proxy in five VPNs: AceVPN, F-Secure’s Freedom VPN, SurfEasy, CyberGhost and VPN Gate. In each case, proxies did not inject additional headers, but consistently modified our existing headers in ways consistent with parsing and subsequent regeneration.

Note that the presence of a proxy does not inherently mean that the VPN is doing something nefarious. VPN Gate’s server software, for instance, appears to systematically funnel requests through a

proxy (though we can find no documentation of this). But because VPN Gate’s server software is provided as an academic project to avoid censorship, we presume that most participating VPNs are not operating endpoints with malicious intent.

6.2.2 Soliciting Providers. In addition to the VPN-based tests, we also emailed approximately 153 VPN providers (not counting those for whom emails bounced, or for whom we could not find a contact point) pretending to be a representative of a company interested in purchasing data about users. The email address used was on a domain dedicated to the ruse, and the email was drafted to look plausible, offering financial incentives in line with numbers offered by a similar company in this space. We only sent one email per provider, and did not follow up with any provider who responded to us. We considered this an ethical component of our research, as answering emails of this type is within the generally accepted employment duties of the person that would receive such an email at each company.

By far, the most common response to our email was a system-generated response informing us that a ticket was opened, which was subsequently closed without comment. For sites that did respond, most were explicit about their lack of interest (e.g. *“Did you even read what our company does? We literally combat this type of stuff”*). Some did promise that our message was “passed on to the proper team for review” and that they would be in touch if warranted. A few providers, however, expressed some tentative interest: one invited us to contact a staff member directly for further discussion, another simply asked for additional details, while a third said that he would *“check [our] website and services [in the] next days and get back to [us] if it triggers [his] interest”*. No provider clearly jumped at the chance to work with a partner, despite realistic, but significant amounts of money offered.

6.3 VPN server infrastructure

When looking at the number of ASNs per VPN provider, IPVanish and CyberGhost advertise vantage points in over 50 different ASNs. However, for the 767 vantage points analyzed, we identify 748 distinct IP addresses associated with 529 different CIDRs. This number implies that VPN providers may share their server infrastructure between them, either due to bi-lateral partnerships or because of outsourcing the infrastructure to the same cloud or hosting provider.

We performed two checks to identify whether multiple companies use similar infrastructure. First, we checked for exact IP address matches across VPN providers. Two providers, Boxpn and Anonine, share four vantage points, each one of them registered by a different company according to WHOIS records. After we relaxed the matching condition at the CIDR-level, we observed that their vantage points were typically found in the same IP blocks: 11 of them, for 16 and 31 vantage points tested, respectively. Unfortunately, we do not have sufficient information to identify whether they are part of the same company, or if they are re-selling infrastructure provided by a third party. The metadata obtained from their websites does not provide sufficient information about whether they are two separate services for the same company. Nevertheless, manual inspection of the listed network of vantage points announced in their websites shows similarities between them, and also with those provided by EasyHideIP, another VPN provider,




Figure 8: Advertised network of vantage points for Anonine, boxpn, and Easy-Hide-IP.

IP Block	ASN (ISO)	VPNs
82.102.27.0/24	9009 (NO)	IPVanish, airVPN, Cyberghost
94.242.192.0/18	5577 (LU)	Acevpn, Cyberghost, Anonine
139.59.0.0/18	14061 (IN)	ra4wvpn, limeVPN, Ironsocket
169.57.0.0/17	36351 (MX)	AceVPN, TunnelBear, Freedom
179.43.128.0/18	51852 (CH)	IPVanish, AceVPN, Anonine, HMA
185.108.128.0/22	30900 (IE)	AceVPN, TunnelBear, Cyberghost
169.57.0.0/17	30900 (IE)	AceVPN, TunnelBear, Freedom
202.176.4.0/24	55720 (MY)	IPVanish, boxpn, Anonine
209.58.176.0/21	59253 (SG)	hideipVPN, vpnland, Cyberghost

Table 5: IP blocks shared by the vantage points of at least three different VPN providers. The country code represents the advertised localization for the vantage point.

as shown in Figure 8. As of this writing, the DNS resolutions for their Argentinian vantage points `ar.boxpnservers.net`, `ar.anonine.net`, and `ar.ehvpn.com` only vary in the final octet: `200.110.156.{183,184,186}`, respectively. The vantage points used for IPVanish and Cyberghost, as well as HideMyAss and Avast, also showed striking similarity at the IP level. In the latter case, this overlap is likely the result of HideMyAss being acquired by Avast Software Inc. in 2015.

In total, we found 40 VPN services with VPN vantage points in the same CIDR block. Many of these IP blocks belonged to well-known hosting providers like Digital Ocean, LeaseWeb and Softlayer. Table 5 lists the blocks that we associated with more than three providers. Such IP blocks are therefore easy to blacklist and block for web services actively discriminating against VPN users.

6.4 Geographic Distribution

VPNs make a variety of claims about where their vantage points are located geographically. They advertise this geographic distribution as a selling point to improve performance, to evade geographic-based blocking on services such as Netflix, and to ‘enhance privacy’. While most providers work hard to provide vantage points in a variety of physical locations, some VPN providers take steps to deceive geo-IP databases about the physical locations of their vantage points. The industry occasionally calls these ‘virtual locations’. From a VPN provider standpoint, such virtual locations reduce latencies for users while still producing the intended outcome of providing a spoofed vantage point in the user-requested location.

Whether or not these virtual locations are a problem depends on the users’ intended use case. For users interested in jurisdictional arbitrage (for illegal file sharing, for instance), they may care deeply about a vantage point actually being located in a particular country. In other cases, users may only care that a vantage point successfully evades geographic blocking. This use case only requires that a service thinks a vantage point is in the country it claims. Unfortunately, a problem arises when non-tech savvy users, especially those potentially vulnerable looking for anonymity, are unable to identify the key differences between each case.

6.4.1 Comparing with Geo-IP Databases. We compared the claimed location of 626 vantage points against the locations determined by three geo-IP databases: IP2Location Lite [14], MaxMind’s GeoLite2 [20], and Google’s location service (as seen from the Google Maps API [9]). The first two services are freely available, and the third can be used in a limited fashion to geolocate a requester’s IP address.

In all cases, there were significant disparities between claimed and IP-geolocated locations, with the greatest differences coming from the database with the expected highest fidelity. Google provided a prediction of location in 541 cases, of which only 377 (70%) agreed with the claimed location. IP2Location provided estimates for 612 vantage points, agreeing in 552 instances (90%). MaxMind similarly provided a location estimate in 612 cases, with 583 cases (95%) agreeing with the VPN service provider’s claimed location. All VPNs were affected with some form of inconsistency between a Geo-IP database and claimed location.

For all databases, about one third of the inconsistencies were the database claiming a vantage point was hosted in the US when the claimed location was elsewhere, with wide varieties of claimed locations. The difficulty of geolocating IP addresses is well documented [8], and errors exist in each dataset.

6.4.2 Identifying ‘Virtual’ Vantage Points. Some VPN providers admit to using ‘virtual’ vantage points, while others attempt to hide it. Our data allows us to identify several previously-unknown instances

of this behavior by identifying co-location of vantage points and discrepancies between location and pings to known hosts.

In total, we identified discrepancies in six of our 62 VPN providers: HideMyAss, Avira, Le VPN, Freedom IP, MyIP.io and VPNUK. An example is Avira’s ‘US’ vantage point in our dataset. Ping times to known hosts in Germany, Luxembourg, and the Netherlands were all less than 9 ms, while pings to known US hosts ranged between 113 and 173 ms. Such ping times indicate that this vantage point is actually in Europe.

Even in the absence of clear geographic anomalies for pings or traceroutes within a single vantage point, we can also identify anomalies by comparing multiple vantage points of the same VPN provider. One example are the five vantage points we have for MyIP.io, including the US, France, Belgium, Germany and Finland. The US and French vantage points reside in the same /24 IP prefix, as do the Belgian, German and Finnish vantage points, but this shared prefix is not by itself a guarantee of co-location. Some VPNs that manage their own IP space heavily subdivide their ranges into very small prefixes across locations (VPN.ht, for instance, appears to split a /24 across at least five distinct locations). Further analysis from ping times and traceroutes, however, suggests that the US and French vantage points are co-located, likely in Montreal, and the other European locations reside together, likely in the UK.

Figure 9 demonstrates co-location of vantage point at three VPN providers. In each case, the graph shows a distribution of ping times to hosts with a known location, ordered from lowest to highest. Though not shown, though ordered by ping time, the same hosts appear in the same order across vantage points. In the case of the European vantage points of MyIP.io, ping times to a reference host varied by less than 1.5ms across different vantage points. A similar process also identifies ‘virtual’ vantage points in FreedomIP as well.


The provider offering the most ‘virtual’ vantage points by far is HideMyAss. An analysis of more than 150 of their endpoints reveals relatively few physical locations. Dozens of locations in North, Central or South America, for instance, appear to be based out of the Seattle area, while dozens more vantage points appear to be based out of Miami, Prague, London and possibly Berlin. HideMyAss specifically advertises ‘virtual’ locations separately from physical servers, but still includes many clearly virtualized vantage points (like North Korea) in their list of physical servers.

6.5 Traffic leakage


Ideally, the traffic forwarded through the VPN tunnel must be opaque to an in-path observer (*e.g.*, Internet service provider, or surveillance agencies). However, VPN providers can introduce misconfigurations and errors that may undermine users’ privacy and security.

The most significant leakage we observed was as a result of tunnel failure. Our test operates by blocking connections between the VPN client and the server and detects tunnel failure leakage by testing whether it is possible to communicate with an outside host. This test underestimates the number of VPN providers with poor failure handling, as the test must guess how long to wait before assuming that the VPN has adapted (or not) to the change in state, but also must keep the total duration as short as possible for practical reasons.


We observed a total 25 of VPN services (58% of applicable services) leaking user traffic when the tunnel failed, including the



(a) Le VPN



(b) MyIP.io



(c) HideMyAss

Figure 9: The distribution of RTT to web hosts, ordered from lowest RTT to highest, for various vantage points of three different VPN providers. Note the strong correlation between some series—these vantage points are co-located despite claims of being in separate countries. The bottom graph plots 148 different series, each corresponding to a distinct HideMyAss vantage point.

well established VPN providers NordVPN, ExpressVPN, TunnelBear, Hotspot Shield and IPVanish. These five VPNs all feature kill-switches in their clients. However, it is either disabled by default, or is designed to only target a chosen application. For instance,

Leakage	VPN Providers
DNS	Freedome VPN, WorldVPN
IPv6	Buffered VPN, BulletVPN, FlyVPN, HideIPVPN Le VPN, LiquidVPN, Private VPN, Zoog VPN Private Tunnel, Seed4.me, VPN.ht, WorldVPN

Table 6: VPN services which leaked DNS and IPv6 traffic from their client software.

the NordVPN macOS client features the option to terminate a selected application on tunnel failure instead of blocking system-wide traffic [24]. We consider disabled kill-switches to be unsafe default behavior, even when intentionally chosen by the VPN provider.

For a privileged attacker such as an ISP or government, tunnel failures are easy to induce by selective blocking or spoofing TCP RSTs. Thus, in countries like China, Russia and Iran, default-off kill-switches actively put users at risk.

For the 43 VPN services which provided their own clients, we also checked for the presence of IPv6 and DNS leakage. (These tests were disabled for our automated testing, as manually controlling OpenVPN requires setting DNS and IPv6 settings manually.) Two VPNs leaked DNS traffic by not modifying the system’s DNS servers. Furthermore, 12 of the VPNs did not block IPv6 traffic when the service did not support IPv6, and hence leaked traffic going to any IPv6-capable hosts. Table 6 enumerates which VPN providers experienced leakage in their product’s default configuration. We observe that leakages were not as common in our commercial VPN client apps as compared to free mobile VPN apps from prior work [13]. For services that relied on third-party OpenVPN clients such as Tunnelblick or Viscosity, few VPN services provided clear instructions to ensure that users’ VPN clients did not leak DNS and IPv6 traffic (as OpenVPN configuration files do not contain the necessary configuration). We found 20 VPNs using such third-party software.

6.6 Peer-to-peer traffic

None of the VPN providers we measured explicitly claimed to operate by routing client traffic out through the connections of other users, and indeed our results bear that out. We found no DNS query traffic from our systems on the Internet-connected interface that was not directly attributable to silent tunnel failures. This result is further not surprising as most VPN services leverage pre-built VPN protocols such as OpenVPN, L2TP, PPTP or IPSec/IKEv2, none of which support routing traffic through clients without considerable additional effort. We leave a more thorough investigation of P2P-based VPN services as future work.

7 RELATED WORK

Several studies have analyzed the privacy risks associated with VPN services, especially those providers offering mobile VPN support. Perta *et al.* [29] manually analyzed the network behavior, infrastructure, and mobile clients of 14 popular VPN services. The study identified instances of developer-induced bugs and misconfigurations that lead to IPv6 and DNS leaks. An extended follow-up study by Ikram *et al.* [13] analyzed more than 200 Android VPN

apps, identifying instances of malware presence, traffic leakages, traffic manipulations and even TLS interception. Ikram’s work revealed that HotspotShield VPN actively injected JavaScript code to redirect users to partner companies [4]. Finally, Zhang *et al.* investigated security vulnerabilities in 84 OpenVPN-based Android applications [44]. Our work confirmed the findings of previous studies in the mobile space, by reporting instances of vulnerable VPN mobile apps due to insecure custom modification and developer-induced misconfigurations.

Vulnerabilities in VPN Services: Appelbaum *et al.* identified security vulnerabilities of commercial and public online VPN services [2]. Al-Fannah studied the privacy risks of the introduction of WebRTC APIs in modern web browsers [1]. This work demonstrated how the WebRTC API can compromise user anonymity by leaking a range of client IP addresses to a visited website, even if a VPN is in use. In our study, we systematically audit whether commercial VPN services present this vulnerability. Our paper, in turn, presents a method to systematically identify and analyze security and privacy aspects of commercial VPN services.

VPN-based Measurements: As VPNs claim to provide access to vantage points remotely, a number of academic efforts leverage VPN services for censorship analysis [6, 19, 22, 35], and network analysis of ISPs [7, 39]. The findings of this project confirm that running measurements over VPN services requires caution due to their frequent traffic manipulation practices.

Open HTTP Proxies: Users also resort to Open HTTP proxies and other censorship resistance systems to anonymize their traffic [16, 42]. One of these services is VPN Gate, an open and free service analyzed by Nobori *et al.* [23]. Sporadic evidence so far has shown that traffic manipulation and monitoring also happens in other types of network relays, including open HTTP proxies [36] and even anonymity networks [5, 43]. The work by Tsirantonakis *et al.* studied header manipulation performed by over 65,000 open HTTP Proxies [36]: 5% of the tested proxies were found to perform malicious or unwanted modification. This included ad injection, redirections, as well as JavaScript injection used for tracking and user fingerprinting purposes.

8 DISCUSSION & CONCLUSIONS

The results of prior work on Android VPNs [13] and open web proxies [36] suggest that the VPN ecosystem might contain a substantial fraction of providers looking to take advantage of their privileged position between the Internet and their customers. In particular, prior work shows that significant numbers of vantage points intercepted and/or manipulated traffic. In that respect, our findings are largely consistent with these adjacent ecosystems. Though we find a smaller number and fewer types of violations, the challenges of large-scale desktop VPN testing limited our study to a smaller number of services. Further, the vast majority of the VPNs we measure are paid services and may therefore have less incentive to further monetize their customers when compared to free services. Finally, many of the VPN services we consider rely on standardized protocols and client software, reducing the opportunity for providers to perform

malicious activity (*e.g.*, TLS interception or surreptitious routing of client traffic through other clients).

While we did observe one isolated instance of traffic manipulation, its sole purpose was to incentivize users to pay for a subscription. This class of traffic manipulation likely provides greater revenue to the VPN provider than simply injecting ads, which was the most prevalent injection activity in prior studies. We note, however, that our ability to identify targeted instances of traffic manipulation across the Web is limited by the scalability problems of our method. Further, VPN providers can passively inspect all unencrypted traffic passing through the VPN, which no measurement will be able to detect; hence, our findings are likely conservative. The most significant, if circumstantial, indication of questionable behavior is the ubiquitous use of affiliate marketing to advertise their services. Because so much of the information online regarding the relative quality of VPN services is dominated by publishers participating in affiliate programs, honest evaluations of these services are hard to come by.

One of the few user-visible ways VPN services have to differentiate themselves to potential customers is through the set of countries in which they provide vantage points. As a result, VPN providers are incentivized to expand this list as much as possible. In some cases, we see clear evidence of providers expanding this list through ‘virtual’ vantage points rather than physical presence in a given country. While ‘virtual’ vantage points are not inherently bad—many users may simply be seeking to evade geo-fencing rather than engage in jurisdictional arbitrage—many services employing virtual vantage points fail to disclose their potential shortcomings to customers.

Finally, we find that VPN providers of all kinds often have poor default configurations, resulting in unintentional data leakage—especially in the case of a tunnel connection failure. Even in 2018, using a VPN ‘safely’ remains a task mostly beyond the amateur user, and no VPN, at least that we were able to find, is perfect.

We intend to publicly release our VPN insights in the form of an informative website `vpnselection.guide`. Data from our evaluations are also available upon request. The VPN test suite is available at `github.com/tahakhan5/vpn_tests`. This tool allows interested individuals to independently evaluate a VPN service, and we hope it will be helpful for the global community in selecting VPNs services which are reliable and best suit their needs.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation through grants CNS-1629973, CNS-1351058, CNS-1564329, CNS-1705050, and the Open Technology Fund Information Controls Fellowship Program. We would like to thank Kyle Aguilar, Reem Hussein, and Mohammad Abdul Salam at UIC for their help in running tests, as well as Joel Rodiel-Lucero and Katie Meyers for painstakingly purchasing VPN accounts. We also would like to thank Muhammad Ikram for his guidance on VPN testing frameworks, Sadia Afroze for advising Taha in the initial stages of this work, and the OTF community for their suggestions, feedback and support. Finally, we would like to thank our shepherd Vijay Erramilli and the anonymous reviewers for their insightful feedback and suggestions.

REFERENCES

- [1] Nasser Mohammed Al-Fannah. One Leak Will Sink A Ship: WebRTC IP Address Leaks. In *Proceedings of the IEEE International Carnahan Conference on Security Technology (ICCSST)*, Madrid, Spain, October 2017.
- [2] Jacob Appelbaum, Marsh Ray, Karl Koscher, and Ian Finder. vpwns: Virtual Pwned Networks. In *USENIX FOCI*, 2012.
- [3] The Best VPN. Are VPNs Legal In Your Country? <https://thebestvpn.com/are-vpns-legal-banned-countries/>, April 2018.
- [4] Center for Democracy & Technology. Complaint, Request for Investigation, Injunction, and Other Relief. AnchorFree, Inc. Hotspot Shield VPN. <https://cdt.org/files/2017/08/FTC-CDT-VPN-complaint-8-7-17.pdf>, April 2017.
- [5] Sambuddho Chakravarty, Georgios Portokalidis, Michalis Polychronakis, and Angelos D Keromytis. Detecting Traffic Snooping in Tor Using Decoys. In *International Workshop on Recent Advances in Intrusion Detection*. Springer, 2011.
- [6] Shinyoung Cho, Rishab Nithyanand, Abbas Razaghpanah, and Phillipa Gill. A Churn for the Better. In *Proc. ACM Int. Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, 2017.
- [7] Taejoong Chung, David Choffnes, and Alan Mislove. Tunneling for Transparency: A Large-Scale Analysis of End-to-End Violations in the Internet. In *Proc. ACM Int. Measurement Conference (IMC)*. ACM, 2016.
- [8] Manaf Gharaibeh, Anant Shah, Bradley Huffaker, Han Zhang, Roya Ensaifi, and Christos Papadopoulos. A Look at Router Geolocation in Public and Commercial Databases. In *Proc. ACM Int. Measurement Conference (IMC)*. ACM, 2017.
- [9] Google. Geocoding API Introduction. <https://developers.google.com/maps/documentation/geocoding/intro>, 2018.
- [10] GovTrack. S.J.Res. 34: A joint resolution providing for congressional disapproval under chapter 8 of title 5, United States Code, of the rule submitted by the Federal Communications Commission relating to 'Protecting the Privacy of Customers of Broadband and Other Telecommunications Services'. <https://www.govtrack.us/congress/bills/115/sjres34/text>, March 2017.
- [11] Lily Hay Newman. If You Want a VPN to Protect Your Privacy, Start Here. *Wired*, March 2017. <https://www.wired.com/2017/03/want-use-vpn-protect-privacy-start/>.
- [12] Lily Hay Newman. The Attack on Global Privacy Leaves Few Places to Turn. *Wired*, August 2017. <https://www.wired.com/story/china-russia-vpn-crackdown/>.
- [13] 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 *Proc. ACM Int. Measurement Conference (IMC)*, 2016.
- [14] IP2Location. Free IP Geolocation Database. <https://lite.ip2location.com/>, May 2018.
- [15] IT Portal. VPN is harming the future of content producers and this will end. <https://www.itportal.com/features/vpn-is-harming-the-future-of-content-producers-and-this-will-end/>, November 2017.
- [16] Sheharbano Khattak, Tariq Elahi, Laurent Simon, Colleen M Swanson, Steven J Murdoch, and Ian Goldberg. SOK: Making Sense of Censorship Resistance Systems. *Proc. Int. Privacy Enhancing Technologies Symposium (PETS)*, 2016.
- [17] Sheharbano Khattak, Mobin Javed, Syed Ali Khayam, Zartash Afzal Uzmi, and Vern Paxson. A Look at the Consequences of Internet Censorship Through an ISP Lens. In *Proc. ACM Int. Measurement Conference (IMC)*. ACM, 2014.
- [18] Nate Lord. The History of Data Breaches. *Digital Guardian*, April 2018. <https://digitalguardian.com/blog/history-data-breaches>.
- [19] Anuradha Mathrani and Massoud Alipour. Website Blocking Across Ten Countries: A Snapshot. In *PACIS*, 2010.
- [20] MaxMind. GeoIP2 Databases. <https://www.maxmind.com/en/geoip2-databases>, May 2018.
- [21] Mozilla Foundation. Public Suffix List. <https://publicsuffix.org/>, May 2018.
- [22] Zubair Nabi. The Anatomy of Web Censorship in Pakistan. In *FOCI*, 2013.
- [23] Daiyuu Nobori and Yasushi Shinjo. VPN Gate: A Volunteer-Organized Public VPN Relay System with Blocking Resistance for Bypassing Government Censorship Firewalls. In *Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2014.
- [24] NordVPN. Automatic Kill Switch. <https://nordvpn.com/features/kill-switch-technique/>, May 2018.
- [25] NordVPN. DNS Leakage test. <https://nordvpn.com/features/dns-leak-test/>, May 2018.
- [26] The One Privacy Site. <https://thatoneprivacysite.net/>, May 2018.
- [27] The Open Technology Fund. <https://www.opentech.fund/>, May 2018.
- [28] OpenVPN. <https://openvpn.net/>, May 2018.
- [29] 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. *Proc. Int. Privacy Enhancing Technologies Symposium (PETS)*, 2015.
- [30] Reddit: VPN. <https://www.reddit.com/r/VPN/>, May 2018.
- [31] Seed4Me The Blog: How to Install VPN Apps from China's App Store? <https://seed4.me/blog/how-to-install-vpn-from-chinas-app-store/>, May 2018.
- [32] Statista. Size of the virtual private network (VPN) market worldwide by type in 2014 and 2019 (in billion U.S. dollars). <https://www.statista.com/statistics/542797/worldwide-virtual-private-network-market-by-type/>, May 2018.
- [33] Ramesh Subramanian. The Growth of Global Internet Censorship and Circumvention: A Survey. *Communications of the International Information Management Association (CIIMA)*, 11(2), 2011.
- [34] Tor Project. <https://www.torproject.org/>, May 2018.
- [35] Michael Carl Tschantz, Sadia Afroz, Vern Paxson, et al. SoK: Towards Grounding Censorship Circumvention in Empiricism. In *Proc. IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2016.
- [36] Giorgos Tsirantonakis, Panagiotis Iliia, Sotiris Ioannidis, Elias Athanasopoulos, and Michalis Polychronakis. A Large-scale Analysis of Content Modification by Open HTTP Proxies. In *Proc. Network and Distributed System Security Symposium (NDSS)*, 2018.
- [37] TunnelBear. TunnelBear Completes Industry-First Consumer VPN Public Security Audit. https://www.tunnelbear.com/blog/tunnelbear_public_security_audit/, Aug 2017.
- [38] Tunnelblick. <https://tunnelblick.net/>, May 2018.
- [39] Gareth Tyson, Shan Huang, Felix Cuadrado, Ignacio Castro, Vasile C Perta, Arjuna Sathiseelan, and Steve Uhlig. Exploring HTTP Header Manipulation In-The-Wild. In *Proc. of the International Web Conference (WWW)*, 2017.
- [40] VPN Mentor. 5 Best VPNs Guaranteed to Beat Netflix's Block in April 2018. <https://www.vpnmentor.com/blog/5-best-vpns-netflix-actually-work/>, April 2018.
- [41] VPNMentor. <https://www.vpnmentor.com/>, May 2018.
- [42] Yuzhi Wang, Ping Ji, Borui Ye, Pengjun Wang, Rong Luo, and Huazhong Yang. GoHop: Personal VPN to Defend from Censorship. In *Proc. Int. Conference on Advanced Communication Technology (ICACT)*. IEEE, 2014.
- [43] Philipp Winter, Richard Köwer, Martin Mulazzani, Markus Huber, Sebastian Schrittwieser, Stefan Lindskog, and Edgar Weippl. Spoiled Onions: Exposing Malicious Tor Exit Relays. In *Proc. Int. Privacy Enhancing Technologies Symposium (PETS)*. Springer, 2014.
- [44] Qi Zhang, Juanru Li, Yuanyuan Zhang, Hui Wang, and Dawu Gu. Oh-Pwn-VPN! Security Analysis of OpenVPN-based Android Apps. In *Proceedings of the International Conference on Cryptology And Network Security (CANS)*, 2017.

A COMPLETE LIST OF VPN SERVICES

Table 7 lists the complete set of VPNs which were evaluated along with their subscription types. For paid VPN services, we had to purchase a subscription. Trial VPNs were specially the ones which had versions with expired after a certain duration. Free VPN subscriptions did not expire, but rather had limited usage features such as vantage points, bandwidth and capped data transfer speeds.

VPN Name	Subscription
AceVPN	Paid
AirVPN	Paid
Anonie	Paid
Avast	Trial
Avira	Trial
Betternet	Free
Boxpn	Paid
Buffered VPN	Paid
BulletVPN	Paid
Celo.net	Trial
CrypticVPN	Paid
CyberGhost	Paid
Encrypt.me	Trial
ExpressVPN	Paid
FinchVPN	Paid
FlowVPN	Trial
FlyVPN	Paid
Freedome VPN	Paid
Freedom IP	Paid
Goose VPN	Paid
GoTrusted VPN	Paid
HideIPVPN	Trial
HideMyAss	Paid
Hotsopt Shield	Paid
IB VPN	Trial
IPVanish	Paid
Ironsocket	Paid
Le VPN	Paid
LimeVPN	Paid
LiquidVPN	Paid
Mullvad	Paid
MyIP.io	Paid
NordVPN	Paid
NVPN	Paid
PrivateVPN	Trial
Private Tunnel	Trial
Private Internet Access	Paid
ProtonVPN	Free
ProxVPN	Free
PureVPN	Paid
RA4W VPN	Paid
SaferVPN	Trial
SecureVPN	Trial
Seed4.me	Trial
Shade YouVPN	Trial
Shellfire	Free
Steganos Online Shield	Trial
SurfEasy	Trial
SwitchVPN	Trial
TorVPN	Trial
Trust.zone	Trial
TunnelBear	Free
VPNBook	Free
VPNUK	Trial
VPNLand	Trial
VPN Gate	Free
VPN Monster	Trial
VPN.ht	Paid
WorldVPN	Trial
Windscribe	Trial
ZenVPN	Trial
Zoog VPN	Free

Table 7: The List of VPN Services Evaluated