From Scarcity to Opportunity: Examining Abuse of the IPv4 Leasing Market

Bernhard Degen*, Ben Du†, Ricky K. P. Mok†, Raffaele Sommese*,
Mattijs Jonker*, Roland van Rijswijk-Deij*, kc claffy†
*University of Twente, Enschede, The Netherlands
†CAIDA/UC San Diego, La Jolla, CA, USA

***Abstract*—Since the exhaustion of unallocated IP addresses at the Internet Assigned Numbers Authority (IANA), a market for IPv4 addresses has emerged. In complement to purchasing address space, leasing IP addresses is becoming increasingly popular. Leasing provides a cost-effective alternative for organizations that seek to scale up without a high upfront investment. However, malicious actors also benefit from leasing as it enables them to rapidly cycle through different addresses, circumventing security measures such as IP blocklisting. We explore the emerging IP leasing market and its implications for Internet security. We examine leasing market data, leveraging blocklists as an indirect measure of involvement in various forms of network abuse. In February 2025, leased prefixes were 2.89× more likely to be flagged by blocklists compared to non-leased prefixes. This result raises questions about whether the IP leasing market should be subject to closer scrutiny.**

I. Introduction

In 2011, the Internet Assigned Numbers Authority (IANA) allocated its final unused /8 IPv4 address block. Although IPv6 was introduced as an alternative more than a decade earlier, IPv4 addresses remain in high demand. Since the last allocation, IPv4 addresses have become increasingly scarce, prompting regional Internet registries (RIRs) to switch to exhaustion management policies and authorize address blocks transfers. With over 6,184 transfers encompassing 30.2 M addresses in 2024 [1], addresses are actively traded. IP brokerage companies (IP brokers in short) act as intermediaries to facilitate this trade. One such broker reported a total of 757 transfers in 2024, totaling a value of US$60.7 M [2]. Lately, their increasing economic value led to the possibility of using IPv4 addresses as collateral for loans [3], with one company issuing US$206 M in notes secured by their IPv4 assets [4].

The increasing value and growing demand for these numeric identifiers put pressure on holders of blocks of IP addresses to use them efficiently. In addition, major cloud providers [5], [6], [7] allow customers to bring their own IP addresses (BYOIP), decoupling the hosting market from the IP address market.

Apart from selling, organizations owning excess address blocks (or *prefixes*) can lease them out to others in need of address space. Leasing offers a more flexible alternative to trading, with leasing terms starting at one month. However, entities prone to blocklisting, such as commercial virtual private network (VPN) providers [8] and bulletproof hosting providers [9], may use IP leasing to their advantage. By frequently rotating IP address blocks — acquiring "clean" blocks and discarding used blocks — they can circumvent address-based filtering [10]. Moreover, leasing reduces Internet transparency by obscuring the actual user of an IP address block [11]. The entity using the addresses can do so without their details being recorded in WHOIS, complicating efforts to respond to abuse complaints and vulnerability notifications.

Despite the increasing importance of leasing, abuse of the IP leasing market has received little attention in research. We aim to address this gap. We establish a methodology to track on-market leases and analyze the types and frequency of abuse associated with active leases. As an indication of abuse, we leverage a comprehensive, labeled collection of IP blocklists covering various types of network abuse, such as distributed denial-of-service (DDoS) attacks or spreading malware.
We make the following contributions:

1) We collect daily leasing market data from two IP brokers over 82 days, examining the relationship between price and abuse.
2) We develop a methodology to estimate the reputation of leased prefixes over time using longitudinal blocklists and find substantial differences between types of abuse.
3) Considering the most recent snapshot of prefixes routed in February 2025, we show that leased prefixes are 2.89× more likely to be abused compared to non-leased prefixes.

To facilitate reproducibility and extension of this work, we publish the datasets supporting our analysis [12].

II. Background

A. IP allocations

In the 1980s, before RIRs were established, IANA allocated IP address blocks directly to organizations wanting to connect to the Internet. Classful routing prevailed at the time, so these *legacy* allocations typically involved address blocks much larger than organizations' actual needs [13]. Approximately 35.9% of the entire IPv4 address space consists of legacy allocations [14]. After RIRs were established, organizations wanting to obtain addresses would instead apply to the RIR serving their region.

A RIR *allocates* addresses to a local Internet registry (LIR) or Internet service provider (ISP), which in turn can *assign* addresses to end-users or networks. Allocated (portable)

978-3-903176-74-4

address blocks can be used independently of any upstream provider, whereas assigned (non-portable) blocks are intended to be used within the infrastructure of the entity that assigned them. The definitions differ slightly by RIR.

Due to the massive growth of the Internet, three RIRs — ARIN, LACNIC, and RIPE NCC — have depleted their address pools. To get a new allocation, members are placed on a waiting list for recovered address blocks. As of February 2025, waiting list times are 498, 615, and 1417 days for RIPE NCC, ARIN, and LACNIC respectively [15], [16], [17]. Apart from the waiting time, there are additional allocation limits. APNIC only allocates a maximum of a /23 and only to new members [18], [19]. AFRINIC and ARIN maintain a maximum of a /22, while RIPE NCC limits allocations to a /24 [20], [16], [15].

Between 2008 and 2017, all RIRs passed policies authorizing transfers of address blocks, effectively making IP addresses an economic good [21]. Address blocks can be transferred (as an *inter-RIR transfer*) between APNIC, ARIN, LACNIC, and RIPE NCC [22]. AFRINIC currently only supports intra-RIR transfers [23]. These policy changes gave rise to an IP transfer market that enables organizations to trade address blocks. However, the property rights of IP addresses are not legally well-defined [24]. Although the prefix can be directly obtained from a seller, the transfer still needs to be processed at the RIRs.

B. IP leasing

IP leasing provides an alternative way for organizations to utilize address blocks without involvement of the RIRs. In this context, a *lessor* (the IP holder) leases address blocks to a *lessee*, which obtains temporary usage rights of the blocks. Lessees can lease prefixes for as little as one month without incurring transfer costs. Since the IP prefix remains registered to the lessor, the lease does not have to be processed at the RIRs.

As it currently stands, there is ambiguity surrounding the legitimacy of IP leasing. None of the five RIRs explicitly allow or disallow address leasing in their policies [23], [25], [26], [27], [28].

C. IP brokers

IP brokers are companies that facilitate trading and leasing of addresses by guiding both parties through the transfer or lease process. They may operate a marketplace where the provider can list its prefix for sale or lease, while the recipient can browse a catalog to find a prefix that meets their needs. We define the following roles:

IP holder A RIR member that has a portable address block allocated to them.

Lessor An IP holder that for a limited duration grants another party usage rights of (part of) its address space.

Lessee An entity that acquires temporary usage rights from the lessor.

IP broker A company that acts as an intermediary between lessor and lessee.

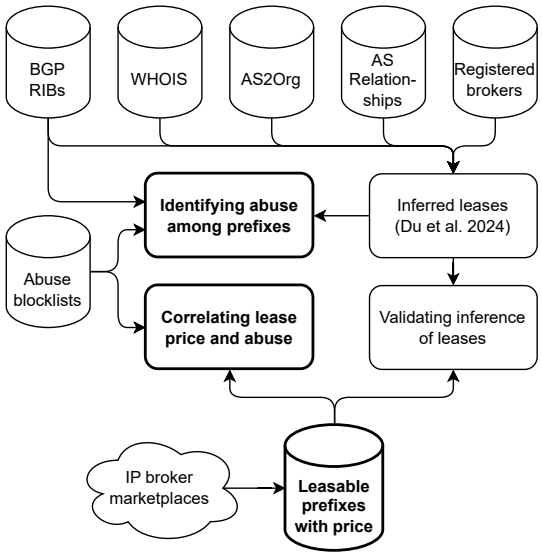

Fig. 1. Overview of data processing. Datasets are represented as cylinders, analysis steps as rectangles, and the brokers' marketplaces as a cloud. The primary outcomes are set in boldface. The top row of datasets serve as input to infer monthly lists of leased prefixes, including the RIR and originating ASN for each lease. We validate the inference step by intersecting the results with prefixes that were available for lease up to a week prior. To investigate whether pricing reflects past abuse, we correlate blocklists with the prices of prefixes available for lease. We use blocklists to determine if and when active leases are blocklisted.

To mitigate the risk of abuse, some IP brokers offer "IP reputation services" to inform their customers of the standing of their prefixes [29], [30], [31]. Several brokers provide instructions on delisting prefixes from blocklists [32], [33], while others provide delisting as a service [34], [35]. Besides reputation management, some IP brokers offer Resource Public Key Infrastructure (RPKI) management [36], [37] or *park* a prefix by announcing it from their autonomous system (AS) [38].

III. Datasets

In this work, we employ the following data sources, as depicted in Figure 1.

IP leasing marketplaces We collect daily marketplace data from two IP brokers between December 21, 2024 and March 12, 2025, adhering to the ethical considerations outlined in Appendix A.

Leasing inference datasets We infer leases using Border Gateway Protocol (BGP) data from RouteViews [39], RIPE RIS [40], monthly WHOIS data, the CAIDA AS2Org [41] and AS Relationships [42] datasets, and the registered broker lists [43], [44]. We have monthly snapshots from September, 2024 up to and including February, 2025.

FireHOL IP lists We collect daily snapshots between January 1, 2024 and March 12, 2025 of the FireHOL IP lists [45], which is a comprehensive composition of multiple blocklists.

BGP routing information base (RIB) We randomly sample routed prefixes of non-leased address space to establish a baseline of abuse for comparing to abuse seen for leased prefixes. We take samples from the multi-hop route-views 2.routeviews.org collector's RIB [39] from February 10, 2025 at 00:00 UTC, the point at which leasing inference of the last snapshot concluded, after removing the leased prefixes.

We make our datasets available to the extent permitted by their respective providers. Specifically, we publish [12] the IP leasing market data for Broker A (with prefixes anonymized), along with the leasing inference data. FireHOL IP lists [45] and the RouteViews RIB [39] are publicly available on their respective websites. Researchers interested in reproducing our results are encouraged to contact the authors.

IV. Methodology

A. Finding IP brokers

As a first step we investigated the IP brokers landscape to obtain marketplace data, and to gain an understanding of how these companies operate. None of the 17 brokers we found provide public marketplace data; however, two brokers provided aggregated market statistics [46], [47]. To check if more information is available to registered members, we attempted to create an account with each broker that allowed registration. One challenge was that most brokers require proof of company ownership as part of their registration process, e.g., ownership of an AS and company registration records.

From the 17 brokers that we found during our non-exhaustive search, we discarded two that only facilitate prefix transfers. We only found an option to sign up for an account with five brokers, and we proceeded to do so using our real names and affiliations. After registration only one broker immediately gave us access to their marketplace. Out of the remaining four, two required us to supply company information, one required us to sign a service agreement (which we declined), and one sent us a personal email message inquiring about "our intentions for buying IPv4 space". After providing company information to the two brokers who requested it, one of them allowed us access to their marketplace.

Thus, we were able to access the marketplaces of two brokers, which we will henceforth refer to as Broker A and Broker B.[1] Broker A is a major player in the market, while Broker B operates on a smaller scale. From December 21, 2024 to March 12, 2025, we collected daily IP leasing market data, capturing all prefixes available for lease along with their monthly leasing price. We verified that the price can be set by the lessor.

We used a publicly documented application programming interface (API) to collect data from Broker A and confirmed with the broker that our access pattern was acceptable. For Broker B, we collected data through their website. To the best of our understanding, this did not violate Broker B's terms of service. Additional ethical considerations are discussed in Appendix A.

IP brokers allow lessors to split prefixes into multiple sizes. Consequentially, the marketplace data contain numerous subdivisions of the same prefix. Therefore, we included only the largest covering prefix for each set of overlapping prefixes. During the measurement period, Broker A and Broker B had a daily average of 393 and 84 non-overlapping prefixes available, respectively (1,771 and 93 without discarding sub-prefixes)

B. Inferring leased prefixes

The data collected from the IP brokers' members interface were limited from a temporal and organizational point of view. To broaden our coverage across different IP brokers and extend the analysis period, we expanded our analysis using the inference method by Du et al. [48] to identify the subset of routed prefixes that are being leased out. In summary:

1) We extract all address blocks from the WHOIS databases and construct a prefix tree from the address blocks. In the resulting tree, allocations made by RIRs are root nodes (portable) and non-portable allocations are leaf nodes.
2) For each root node, we find all ASes assigned to the organizations of the root node.
3) For each leaf node, we find the most specific covering prefix and origin ASes from BGP.
4) We infer a prefix as *leased* if the leaf node is originated exclusively by ASes that have no known AS relationship with: (i) any of the ASes originating the root prefix (if it has a BGP origin); or (ii) any of the ASes assigned to the root node (if the root does not have a BGP origin, step 2).

To match the monthly granularity of WHOIS data, we generate monthly snapshots of inferred leases. Each snapshot contains the leased prefixes and the ASes announcing them up to the 10th day of that month; any prefixes announced later appear in the snapshot of the next month.

We note several limitations of this inference method. It may incorrectly identify certain prefixes as leased, including squatted and hijacked prefixes, as well as those originated by BGP-based DDoS-scrubbing providers or parked by IP brokers on behalf of their customers. Further, as a consequence of step 3, we can only infer leased prefixes that are announced in BGP. Hence, we do not have visibility into prefixes that are leased but unused. Finally, the method does not distinguish between a prefix announced by multiple ASes during disjoint time intervals and a multiple origin AS (MOAS). As a consequence, the method may incorrectly infer these prefixes as leased.

Lessees may choose to deaggregate leased prefixes and announce the resulting subprefixes individually. Since we lack direct information about the exact prefix size agreed upon between lessor and lessee, we group contiguous prefixes announced by the same origin. We observe that a notable fraction of these groups (20.1—20.5%) in the leasing inference dataset contain multiple contiguous prefixes announced separately.

Du et al. [48] validated their inferences against prefixes managed by organizations on the registered brokers list,

[1] We anonymize the brokers to prevent inadvertent reputational damage.

achieving a precision (i.e., the ratio of correctly inferred leases to total inferred leases) of 98% and recall (i.e., the ratio of correctly inferred leases to total leases) of 82%. As this dataset is central to our study, we perform additional validation to verify that the inference method can identify prefixes offered on leasing marketplaces for which we collected data in §IV-A. In particular, we selected the inference period that overlaps with the period for which we have marketplace data, from January 11, 2025 through February 10, 2025. We selected prefixes of any length that were removed from marketplaces in the week before the inference period and that did not reappear on these marketplaces during the inference period. Next, we filtered for these prefixes in the set of inferred prefixes, either announced using the same prefix length or in aggregate. Of the 366 prefixes that disappeared from the marketplace, we inferred 129 (35.2%) as leased. Although the intersection is incomplete, we note that the inference method detected only announced prefixes and that we cannot confirm if prefixes removed from the marketplace were successfully leased.

C. Using blocklists to infer abuse

Blocklists are lists of Internet resources identified as being involved in malicious activities, which operators use to protect their networks and researchers use to characterize abuse [49], [10]. Blocklisted resources typically include ASNs, IP addresses, prefixes, domain names, or a combination. For our study, we used blocklists containing exclusively IPv4 addresses and prefixes, as these resources most closely align with the subject of our study. IP blocklists more accurately reflect malicious use of specific IP addresses, whereas domain name or ASN blocklists may indicate indirect abuse not necessarily attributable to the users of the IP addresses themselves. For example, an ASN may originate prefixes with mixed reputations. Additionally, inaccuracies may arise from the additional translation steps required when using ASNs or domain name blocklists. For this reason, we decided to use daily snapshots of the FireHOL IP lists [45] between January 1, 2024 and March 12, 2025. We consider a prefix blocklisted if at least one of its IP addresses matches an address or overlaps with a prefix on the blocklists.

FireHOL aggregates 253 blocklists and groups them each into: `abuse`, `anonymizers`, `attacks`, `malware`, `reputation`, `organizations`, `spam`, `unroutable`. These categories are defined as follows: The `abuse` category monitors automated web scripts (bots) that perform actions such as leaving unwanted comments or attempting to log in to websites. The `attacks` category covers various forms of hacking and online threats. The `anonymizers` category lists open proxies and Tor exit nodes. The `malware` category tracks worms, command and control (C2) servers, and addresses distributing malware. The `organizations` category primarily consists of IP ranges belonging to legitimate organizations, such as ISPs and software companies. The `reputations` category tracks addresses linked to criminal operations and sharing of harmful content. The `spam` category focuses on email spam. The `unroutable` category tracks bogons: unallocated or reserved IP addresses.

Table I breaks down the included lists by category, aligned with the start date of the leasing market dataset. It shows per category the number of non-empty lists, entries (either an IP address or prefix), and the equivalent number of addresses. As each category aggregates blocklists from various sources, entries and addresses may be duplicated. We report the number of unique (possibly overlapping) entries and the number of unique addresses within each category. We observed that those blocklist entries are heavily skewed towards long prefixes, with individual IP addresses accounting for 97.0% of all entries. Only 0.248% of the entries in the blocklists are prefixes shorter than /24, indicating targeted blocking of small subnets.

We decided to exclude `organizations` and `unroutable` categories from our analysis, since the former is not indicative of malicious behavior, and the latter cannot be used on the Internet. Although `anonymizers` do not inherently indicate abuse, such services can enable it; therefore, we have chosen to include this category. Furthermore, we found that the `attacks` category included bogon ranges (those in the `unroutable` category), accounting for 588,513,773 IP addresses (86.2% of all addresses within the category). As this artificially inflates the number of represented addresses, we remove bogon ranges from the statistics presented in this section. We verified that no prefixes in the marketplace dataset or the inferred leases dataset were bogons; therefore, this does not impact our subsequent analysis.

Another property of a blocklist is the retention period, which is how long an entry remains on the list. Several lists have a retention period of less than a day or list every entry only once. To avoid missing these short-lived blocklist entries, we included the most recent daily aggregated versions (suffixed `_1d`) of these lists. Although we do not control the intervals or offsets at which these lists are aggregated, fetching all lists every 24 hours ensures that all entries are eventually captured.

We observe substantial differences between categories. `anonymizers` mainly consisted of individual addresses, whereas `malware` included only 5.8% of the entries but covered 62.0% of all blocklisted addresses. The `spam` category contained an order of magnitude fewer addresses than the next smallest category. Future work should include an in-depth analysis of blocklists to account for and normalize differences in their scope, selection criteria, and update frequency. This would support more robust and nuanced comparisons of the specific types of abuse observed in prefixes.

In addition to the aggregated 253 blocklists, the FireHOL authors provide a much shorter *Level 1* blocklist that is considered safe to use on firewalls as it strives for zero false positives [45]. The Level 1 blocklist combines eight lists in the `attacks`, `malware`, and `reputation` categories that flag addresses with high confidence. On December 21, 2024, this blocklist contained 4,191 unique, non-bogon entries, representing 24,370,176 addresses. Prefixes inferred leased between December 11, 2024 and January 10, 2025 cover 4.20%

TABLE I
FIREHOL BLOCKLIST STATISTICS ON DECEMBER 21, 2024

Category	Lists		Entries				Addresses			
	N	%	*N*	%	Unique	%	*N*	%	Unique	%
`abuse`	21	10.5	291,064	14.3	165,081	9.8	740,481	0.4	462,439	0.9
`anonymizers`	19	9.5	1,135,277	55.8	1,095,388	65.3	1,270,082	0.6	988,347	1.8
`attacks`	56	28.0	279,902	13.8	149,264	8.9	94,199,004	45.1	18,008,687	33.5
`malware`	67	33.5	118,349	5.8	101,198	6.0	35,535,960	17.0	33,337,740	62.0
`reputation`	20	10.0	138,669	6.8	102,000	6.1	76,883,141	36.8	898,257	1.7
`spam`	17	8.5	71,089	3.5	65,048	3.9	71,088	0.0	58,544	0.1

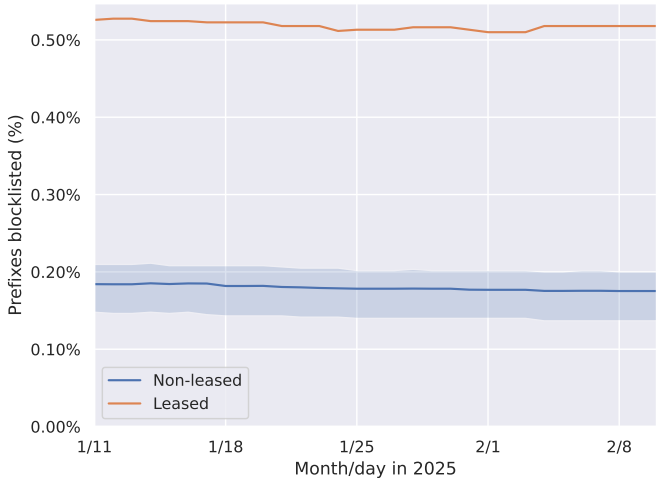

Fig. 2. Fraction of leased and non-leased prefixes flagged by the FireHOL Level 1 blocklist between January 11 and February 10, 2025. The shaded region indicates the range of values observed in the 10 samples of non-leased prefixes.

and 13.5% of the entries on the Level 1 and full aggregated blocklists, respectively. When comparing the different types of abuse, we use the full aggregated list of 253 blocklists, as it has labeled categories (excluding the `organizations` and `unroutable` entries). However, when we compare blocklisting trends in different groups, e.g., leased vs. non-leased prefixes, we use the Level 1 blocklist, which provides higher confidence.

V. RESULTS

A. Prevalence of abuse in leased prefixes

To evaluate whether leased prefixes are associated with increased levels of abuse, we compared the blocklisting prevalence of leased prefixes to a baseline of non-leased routed prefixes. We started with the most recent February snapshot of inferred leases, containing leased prefixes between January 11 and February 10, 2025. We then collected all routed prefixes from the BGP RIB [39] of February 10 at midnight UTC, corresponding to the time the leased prefix inference was concluded. From this set of routed prefixes, we removed those that overlapped with any leased prefix in the snapshot, resulting in a set of non-leased routed prefixes. To ensure a fair comparison, we randomly sampled a non-leased prefix with the same prefix length for each leased prefix, resulting in two groups of equal size. A limitation is that the selected non-leased IP space could include more routed but unused ranges, potentially underestimating abuse rates, while leased space may have higher utilization due to associated recurring costs.

For each day of the month, we calculated the ratio of prefixes flagged by the FireHOL Level 1 blocklist in both groups. The FireHOL Level 1 blocklist combines eight lists in the `attacks`, `malware`, and `reputation` categories, which we use to compare leased and non-leased prefixes.

We inferred that from January 11 to February 10, 2025, 62,944 of the 1,032,086 (6.10%) routed prefixes were leased. Figure 2 shows a stable trend in the fraction of blocklisted prefixes during this period. To ensure our sample is representative, we repeated the sampling process of non-leased prefixes 10 times. For each sample, we computed daily blocklisting rates and averaged them over the month. Across these monthly averages, non-leased prefixes were blocklisted at a mean rate of 0.179%, while leased prefixes appeared on blocklists at a mean rate of 0.518%. That is, leased prefixes were 2.89× more likely to be blocklisted in this dataset.

For completeness, we repeated the analysis using the full aggregated blocklist, but excluding the categories `organizations` and `unroutable`. During the same time frame, 17.4% of non-leased and 32.0% of leased prefixes appeared on this longer blocklist. Using this list, leased prefixes were only 1.84× more likely to be blocklisted than non-leased prefixes. Since the inclusion criteria for the Level 1 blocklist are rather conservative, we consider the result a lower bound on the actual number of abused prefixes. Conversely, the complete aggregated blocklist contains a wide variety of blocklists, so we consider the result an upper bound.

B. Types of abuse

We analyzed the types of abuse associated with leasing by comparing leased and non-leased routed prefixes in the most recent February snapshot of inferred leases. For each of the 62,944 leased prefixes, we randomly sampled a non-leased prefix of the same size in the same way as described in §V-A. We then determined the proportion of these prefixes that appeared on the aggregated blocklist for each category indicating abuse during the leasing inference period, i.e., from January 11 through February 10. Figure 3 plots the distribution of these

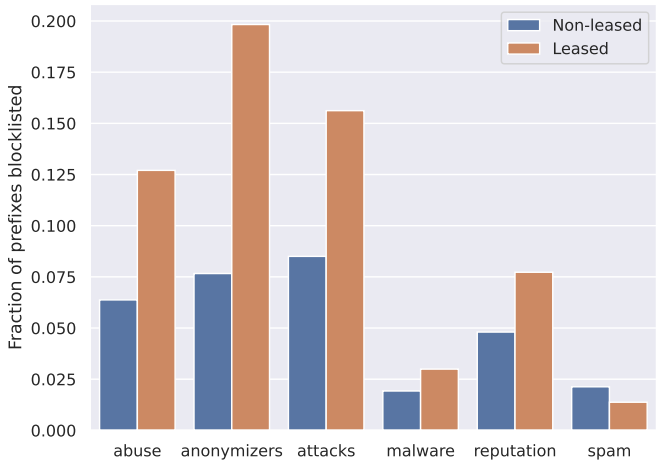

Fig. 3. Fraction of leased and non-leased prefixes flagged by blocklist category (x-axis) between January 11 and February 10, 2025. Blocklisting was more prevalent for leased prefixes in all categories except spam. Some prefixes were flagged in multiple categories.

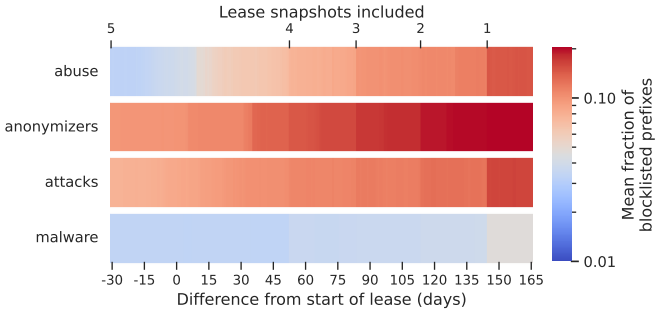

Fig. 4. Prefix reputation after leasing started. The plot shows the mean fraction of blocklisted prefixes that we inferred to be newly leased from September 11, 2024, to February 10, 2025. For brevity, we omit categories exhibiting no discernible change. The bottom x-axis shows the offset t days since the start of the lease, while the top x-axis shows the number of leased snapshots included in computation of the mean. The y-axis indicates the blocklist category.

proportions. Prefixes may appear in multiple categories, but each can appear no more than once in a single category. To ensure our random sample was representative, we sampled 10 times but did not observe any deviations within a category larger than 0.003.

Leased prefixes dominated all but one type of abuse. The difference was most pronounced for the `anonymizers` category, where leased prefixes were 2.59× more prevalent than non-leased prefixes. Interestingly, a leased prefix was 1.55× less likely to be flagged for sending spam. However, comparisons between categories are risky due to the considerable differences in the number, sizes, and composition of blocklists in each category (§IV-C).

C. Temporal dynamics

We then investigated evidence of abuse emerging after a lease starts. For this analysis, we used the monthly snapshots of leased prefixes (§IV-B), and daily blocklists as a proxy for various types of abuse (§IV-C).

We calculated blocklisting activity relative to the start of a lease as follows. The monthly snapshots of inferred leases spanned from September 11, 2024, through February 10, 2025. We aligned the lease start date $t = 0$ to the 26$^{\text{th}}$ day of every month, the approximate midpoint between two subsequent snapshots. Given the monthly granularity of inferring leases, the estimated start dates have an error margin of ± 16 days.

Each lease includes the ASNs originating the leased prefix. We used this identifier as an indication of different leases, i.e., the lessor changing. We discarded inferred monthly leases of prefixes with multiple origins (3.58%), as we could not determine whether the ASNs belonged to one or several lessees.

Next, for each prefix, we determined the date it was first leased by an AS (from September 26 onward). These steps resulted in 18,758 leases associated with 17,196 prefixes originated by 2,661 ASNs. Finally, we determined the daily fraction of prefixes flagged by at least one blocklist within each category. We calculated the offset of days since the start of the lease, and then took the mean value of these fractions.

A limitation of this approach is that the calculated reputation of prefixes is less accurate as time progresses further from the estimated beginning of a lease. The reason is that, at the time of the analysis, not all snapshots started 165 days in the past. We indicate the confidence of our calculation by the number of lease snapshots over which we calculated the mean.

Figure 4 shows the resulting visualization of prefix reputation over time. The fraction of blocklisted prefixes was relatively stable before the start of the lease, but increased steadily after the start of the lease and continued to rise until the end of our measurements. Note that before the start of the lease, the prefixes may not have been in active use and therefore may appear more legitimate than used prefixes. Refer to §V-A for a comparison to routed prefixes.

When comparing the blocklist activity at the start of the lease to 30 days after — one month being the minimal lease duration — we observed a negligible increase for `malware` but an increase of 60.2% in blocklisting rate for `abuse`. Considering instead 150 days after the start of the lease, we observed an increase ranging from 37.9% in blocklisting rate for `malware` to 269% for `abuse`.

We found substantial variation between blocklist categories, e.g., strong increases in the `abuse`, `anonymizers`, and `attacks` categories, but no discernible changes in the `spam` and `reputation` categories, which we exclude from the graph for clarity. For the `malware` category, we ascribe the effect to leases started in the October snapshot, thus we hesitate to generalize this result. We did not observe blocklisting decreasing in any category. Notably, the high prevalence of anonymizing services within leased IP space may attract or facilitate abuse that is later reflected in other categories.

In summary, while for some categories blocklisting increases steadily after the start of a lease, other blocklisting

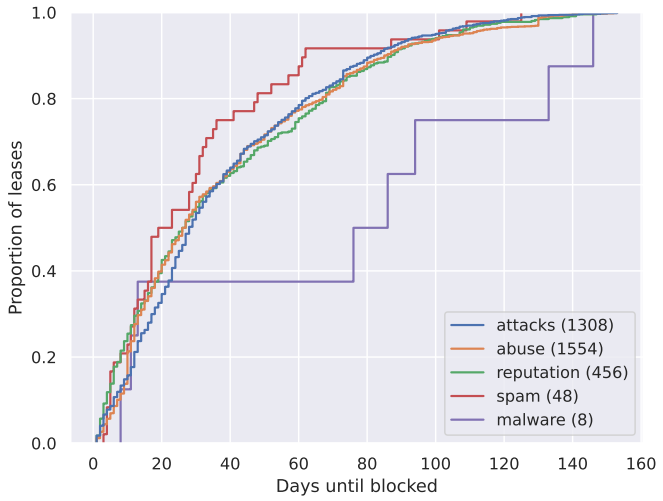

Fig. 5. eCDF of the time until a prefix is blocked. The plot shows that blocklisting activity decreases over time. The number of leases per category is given in parentheses. We exclude the `anonymizers` category as a few large monthly updates distort the plot.

categories show little change. Although lessors have the ability to rotate prefixes after one month, the fraction of blocklisted prefixes continues to rise even five months after a lease begins. This behavior suggests that malicious actors are not frequently rotating their leased IP addresses.

We also analyzed the time elapsed between a prefix being leased and its appearance on abuse-related blocklists. This time span provides an upper bound on how quickly a prefix is abused, as blocklisting requires additional time.

We first identified continuous leases from monthly lease snapshots from September through February. As before, we set the start date to the approximate midpoint between two subsequent snapshots; the 26$^{\text{th}}$ of the month. We determined each lease duration by finding the first and last month its prefix was subsequently originated by the same ASN (excluding those with multiple ASNs).

We discarded leases active in the first month; in the absence of prior data, we cannot determine if these are new or continued leases. Including them would overestimate blocklist timing, which we would calculate relative to the start of the first snapshot instead of to the start of the lease. However, we kept 6.02% of the leases ongoing in the last month, which risks missing those that were blocklisted after the end of the final snapshot. Essentially, we trade off accuracy near the end of leases for data volume.

For each category and lease, we calculated the time offset between the start date and the first time the leased prefix was blocklisted. Due to the monthly inference granularity, we have an error margin of ± 16 days. Since we are interested in prefixes that became blocklisted during their lease, we discarded all leases that were already blocklisted at the start of the lease, as well as those that were never blocklisted during their lease period. Since the `anonymizers` category was dominated by few infrequent, disproportionate updates, we excluded it from this analysis. Finally, we joined the leases with the contemporaneous aggregated blocklist.

The eCDF in Figure 5 shows the distribution of the time until leased prefixes are blocklisted. We observe that most abuse is detected relatively fast, with a median detection time of 27 days across all categories. When comparing categories, we found minimal differences. We included the `spam` and `malware` categories for completeness, but with only 48 and 8 samples respectively, they are too small to support reliable analysis or to draw meaningful conclusions about detecting these types of abuse.

Finally, we examined the relationship between lease duration and blocklisting, i.e., were shorter leases more likely abused? We began by identifying all continued leases from September 11, 2024, through February 10, 2025, excluding those first observed in September. We grouped the 20,351 resulting leases by their duration in months and found that 44.4% lasted only one month, while 9.57% were active for 5 months or longer. As noted in §IV-B, short-lived events such as BGP parking may be misclassified as leases, thereby inflating the number of apparent short-duration leases.

Next, we evaluated how many of these leases were flagged by blocklists at any point during their lifespan. Since the exact start and end date of a lease are unknown, we assumed and allowed the widest possible date range: considering a lease to have started on the 11$^{\text{th}}$ of the first month and ended on the 10$^{\text{th}}$ of the last month. Among these leases, we observed only 136 (0.668%) as blocklisted at any point during their lease.

We grouped the leases based on whether they were blocklisted and calculated the fraction of leases within each group according to their duration in months. We observed similar blocklisting rates across all duration groups, except for the longest group of 5 months and longer, where the ratio of blocklisted prefixes was 1.85$\times$ higher than that of non-blocklisted ones, suggesting that leases continued to be flagged even five months after the lease started. Unfortunately, the overall number of blocklisted leases was substantially smaller than non-blocklisted ones, which limits the generalizability of this finding. Lease duration trends across blocklist categories may offer additional insights and merit future investigation.

D. Lease availability and pricing

We also investigated the relationship between the ask price and reputation of a prefix. We collected marketplace data between December 21, 2024 and March 12, 2025. The data consist of ask prices for prefixes available for a lease of one month. Figure 6 illustrates the volume and median leasing price for prefixes of various sizes on the two brokers we analyzed.

During the observation period, the median monthly leasing price for an address through Broker A was US$0.53 and US$0.47 through Broker B. This 12.8% difference may stem from differences in the fee calculation on both platforms or a difference in features. The contemporaneous median trading price per address was US$33.93. In comparison, with

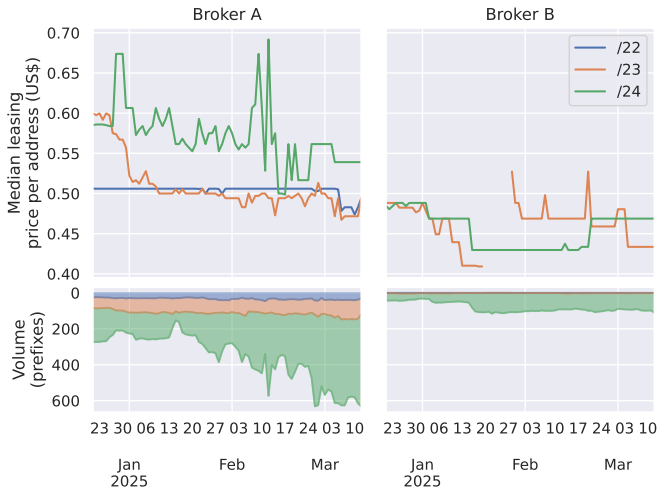

Fig. 6. Development of monthly leasing ask prices. We omit prefixes shorter than /22 for Broker A and shorter than /23 for Broker B from our analysis, since we do not have sufficient data points. Between January 20–27, no /23 prefixes were offered through Broker B.

a monthly leasing price of US$0.50, the amortization period of buying an address block is 68 months or 5.7 years.

We examined the relationship between leasing price and reputation by combining marketplace data with our aggregated blocklist. Specifically, for each prefix, we calculated the median price for each day we observed it on the marketplace within a given a month, and counted the number of matching blocklist entries that expired at most 30 days before the prefix was first listed for lease. We then computed the Pearson correlation coefficient between the one-month leasing prefix price of each prefix and the number of entries on blocklists that flagged it.

We used two full months of data, January and February, 2025. To mitigate differences caused by price variations across prefix sizes and time, we included only the most prevalent prefix size (/24). Additionally, this prefix size is the most specific universally routable size, which reduces variability resulting from concentrated abuse within a large subnet. Due to insufficient data from Broker B (165 for January, and 177 for February), we excluded it from our analysis. In February, we observed $N = 930$ non-overlapping prefixes. We did not find any statistically significant correlation ($r = -0.0070$, $p = 0.8321$) between prefix price and the number of blocklist entries. We repeated this experiment for January during which fewer prefixes were available ($N = 467$) and still found no correlation ($r = -0.0052$, $p = 0.9108$).

In summary, our analysis did not detect a significant relationship between prefix price and past abuse, suggesting that other factors dominate pricing decisions. One possible explanation is suboptimal decision-making due to a lack of market transparency. There is information asymmetry between the lessors and lessees, as lessors may not disclose how a prefix was abused in the past. Conversely, lessors might be unaware of abuse that occurs during a lease. As a result, prefix pricing may be arbitrary and not reflect reputation. Beyond market inefficiencies, blocklist entries may indicate dangling domains in broken links that remain embedded on websites or emails, generating residual traffic that may be valuable to potential lessees. Previous research has shown that the majority of revenue from email spam campaigns is generated *after* the associated domain has been blocklisted [50]. We leave further analysis of leasing market dynamics for future studies.

VI. Related work

Analyzing the IPv4 transfer market: Mueller and Kuerbis [51] analyzed the effect of policy changes passed by RIRs that authorize IPv4 address transfers. They found that 88% of the transferred addresses came from legacy address space and that, unlike the organizations releasing the address blocks, the organizations receiving them were almost exclusively connectivity and online service providers. Given these observations, the authors concluded that the transfer market was effectively redistributing inefficiently allocated legacy blocks. Livadariu et al. [52], [21] used lists of transfers published by RIRs to study the transfer market from November 2009 to December 2015. They compared the median fraction of utilized address space in transferred prefixes 12 months before and after transfer and found an increase of at least 50% across all RIRs. This suggests that buyers acquired addresses for actual operational needs rather than for speculative investment.

Inferring leased address space: Unlike IP transfers, IP leases are not registered with RIRs, requiring inference from external data sources. Prehn et al. [53] proposed inferring leased prefixes by identifying BGP *delegations* — subprefixes originated by an AS different from the one originating the most-specific covering prefix. Du et al. [48] extended their methodology by integrating WHOIS and AS relationships data to exclude prefixes where no business relationship exists between a subprefix and covering prefix and found that 4.1% of all routed prefixes were leased in April 2024. Our work applied the methodology of Du et al. to generate a list of leased prefixes for our abuse investigation (§IV-B).

Identifying malicious activity: Giotsas et al. [10] estimated the abuse of IPv4 prefixes using longitudinal blocklists. They found that, between October 2009 and August 2019, transferred prefixes were overrepresented on blocklists by a factor of 4× to 43× compared to non-transferred prefixes. Unlike their focus on transfers, we examine the leasing market, distinguishing two phases: when a prefix was available and when it was actively leased. Du et al. [48] compared blocklisting of ASes originating leased prefixes to those originating non-leased prefixes and found that the former are about 5× more likely to be blocklisted. However, their reliance on a single, AS-level blocklist offered limited coverage. We used a labeled collection of IP blocklists, providing insight into the misuse of leased prefixes. Aside from leasing address blocks, individual IP addresses can also be leased for use as residential proxies. In this model, bandwidth brokers enable users to monetize access to their home Internet connections. Khan et al. [54] investigated how residential proxies can facilitate fraudulent

activities, and uncovered signs that residential addresses are being used in phishing campaigns.

VII. Discussion and Conclusion

Although the IP leasing market is still young, our findings suggest that its various stakeholders can benefit from greater transparency and oversight. In our analysis of leasing and blocklist datasets collected, we found that leased prefixes were 2.89× more likely to appear on blocklists compared to non-leased prefixes, and that they were disproportionately flagged for all types of malicious activities examined, except for sending spam.

Based on our findings, we provide recommendations for parties involved in IP leasing. Lessors could track the reputation of their prefixes throughout the lease. Lessees could independently assess the reputation of a prefix before leasing it, and we recommend against using the lease price as an indicator of reputation.

IP brokers can also leverage our findings to enhance operational security on their platforms. To protect lessors and lessees, IP brokers can implement sufficient monitoring and enable swift responses to potential abuse. Furthermore, we argue that reputation protection could be a value-added service that IP brokers can offer to lessors as insurance against abuse [55], [56].

Blocklist providers and security specialists can use our insights to identify evasion attempts that exploit IP leasing. Our work builds on prior research in blocklist compilation and demonstrates how such data can aid in classifying and mitigating abuse.

A comprehensive view of the leasing landscape can also assist regulators in evaluating whether the leasing market has a net positive or negative impact. This understanding may inform RIR policies on IP leasing, shaping the market's future.

Future work should focus on performing a detailed residual traffic analysis of leased prefixes after the expiration of the lease to identify trends and anomalies. Another interesting direction is to compare leased and non-leased prefixes belonging to the same IP holder, which could reveal differences in how leased IP space is managed. These investigations could provide further insight into the potential abuse of leased prefixes to lessors and IP brokers, leading to a cleaner Internet.

Acknowledgements

We thank our shepherd and the anonymous reviewers for their insightful comments. We also extend our gratitude to the teams at the IP brokers in this study for their support and for reviewing a pre-publication version of this paper. This research received operational support from the Twente University Centre for Cybersecurity Research (TUCCR).

This material is based on research sponsored by the National Science Foundation (NSF) grants NSF CNS-2120399 and OAC-2131987, and by the European Commission under the GN5-2 project (grant agreement ID 101194278). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the funding agencies.

References

[1] G. Huston, "ISP Column," Jan. 2025. [Online]. Available: https://www.potaroo.net/ispcol/2025-01/addr2024.html

[2] IPv4.Global, "IPv4 Prior sales," Jan. 2025. [Online]. Available: https://auctions.ipv4.global/prior-sales

[3] ——, "IPv4.Global Launches New Lending Program Using IP Addresses as Collateral," Mar. 2025. [Online]. Available: https://circleid.com/posts/ipv4.global-launches-new-lending-program-using-ip-addresses-as-collateral

[4] Cogent, "Cogent Announces IPv4 Address Securitization Offering," Apr. 2024. [Online]. Available: https://www.cogentco.com/en/about-cogent/press-releases/4443-cogent-announces-ipv4-address-securitization-offering

[5] Amazon, "Bring your own IP addresses (BYOIP) to Amazon EC2 - Amazon Elastic Compute Cloud," 2025. [Online]. Available: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-byoip.html

[6] Google, "Bring your own IP addresses | VPC," 2025. [Online]. Available: https://cloud.google.com/vpc/docs/bring-your-own-ip

[7] Microsoft, "Custom IP address prefix (BYOIP) - Azure Virtual Network," Oct. 2024. [Online]. Available: https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/custom-ip-address-prefix

[8] E. Khan, A. Sperotto, J. van der Ham, and R. van Rijswijk-Deij, "Stranger VPNs: Investigating the Geo-Unblocking Capabilities of Commercial VPN Providers," in *Passive and Active Measurement*, A. Brunstrom, M. Flores, and M. Fiore, Eds. Cham: Springer Nature Switzerland, 2023, pp. 46–68.

[9] S. Alrwais, X. Liao, X. Mi, P. Wang, X. Wang, F. Qian, R. Beyah, and D. McCoy, "Under the Shadow of Sunshine: Understanding and Detecting Bulletproof Hosting on Legitimate Service Provider Networks," in *2017 IEEE Symposium on Security and Privacy (SP)*, May 2017, pp. 805–823. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7958611

[10] V. Giotsas, I. Livadariu, and P. Gigis, "A First Look at the Misuse and Abuse of the IPv4 Transfer Market," in *Passive and Active Measurement*, A. Sperotto, A. Dainotti, and B. Stiller, Eds., vol. 12048. Cham: Springer International Publishing, 2020, pp. 88–103. [Online]. Available: http://link.springer.com/10.1007/978-3-030-44081-7_6

[11] B. Kuerbis, "Can't sell your IPv4 numbers? Try leasing them." Apr. 2013. [Online]. Available: https://www.internetgovernance.org/2013/04/28/cant-sell-your-ipv4-numbers-try-leasing-them/

[12] B. Degen, B. Du, R. K. P. Mok, R. Sommese, M. Jonker, R. van Rijswijk-Deij, and k. claffy, "From Scarcity to Opportunity: Examining Abuse of the IPv4 Leasing Market," 2025. [Online]. Available: https://doi.org/10.5281/zenodo.15398237

[13] V. Fuller, T. Li, J. Yu, and K. Varadhan, "Classless inter-domain routing (CIDR): An address assignment and aggregation strategy," IETF, RFC 1519, Sep. 1993. [Online]. Available: http://tools.ietf.org/rfc/rfc1519.txt

[14] Internet Assigned Numbers Authority (IANA), "IANA IPv4 Address Space Registry," Dec. 2023. [Online]. Available: https://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml

[15] RIPE Network Coordination Centre, "IPv4 Waiting List," Feb. 2025. [Online]. Available: https://www.ripe.net/manage-ips-and-asns/ipv4/ipv4-waiting-list/

[16] ARIN, "IPv4 Waiting List," Feb. 2025. [Online]. Available: https://www.arin.net/resources/guide/ipv4/waiting_list/

[17] LACNIC, "IPv4 Address Waitlist," Feb. 2025. [Online]. Available: https://www.lacnic.net/6335/2/lacnic/ipv4-address-waitlist

[18] APNIC, "IPv4 exhaustion | APNIC," 2025. [Online]. Available: https://www.apnic.net/manage-ip/ipv4-exhaustion/

[19] Aftab Siddiqui, "Prop-129: Abolish Waiting list for unmet IPv4 requests | APNIC," Jul. 2019. [Online]. Available: https://www.apnic.net/community/policy/proposals/prop-129

[20] AFRINIC, "AFRINIC IPv4 Exhaustion," Jan. 2023. [Online]. Available: https://afrinic.net/exhaustion

[21] I. Livadariu, A. Elmokashfi, and A. Dhamdhere, "On IPv4 transfer markets: Analyzing reported transfers and inferring transfers in the wild," *Computer Communications*, vol. 111, pp. 105–119, Oct. 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0140366417301305

[22] RIPE Network Coordination Centre, "Inter-RIR Transfers," 2025. [Online]. Available: https://www.ripe.net/manage-ips-and-asns/resource-transfers-and-mergers/inter-rir-transfers/
[23] AFRINIC, "Consolidated Policy Manual (v1.6 Current)," Nov. 2020. [Online]. Available: https://afrinic.net/policy/manual
[24] B. Shantz, "Determining Ownership and Control of IPv4 Addresses Notes," *Washington University Law Review*, vol. 94, no. 3, pp. 739–772, 2017. [Online]. Available: https://heinonline.org/HOL/P?h=hein.journals/walq94&i=757
[25] APNIC, "APNIC Internet Number Resource Policies," Mar. 2015. [Online]. Available: https://www.apnic.net/community/policy/resources
[26] ARIN, "Number Resource Policy Manual," Mar. 2025. [Online]. Available: https://www.arin.net/participate/policy/nrpm/
[27] LACNIC, "LACNIC Policy Manual," LACNIC, Tech. Rep., Jul. 2024. [Online]. Available: https://www.lacnic.net/innovaportal/file/680/1/manual-politicas-en-2-20.pdf
[28] RIPE Network Coordination Centre, "RIPE Policies," 2025. [Online]. Available: https://www.ripe.net/publications/docs/ripe-policies/
[29] IPXO, "Lease IPv4 from IPXO Marketplace. Instant Setup with LOA or ROA," 2025. [Online]. Available: https://www.ipxo.com/lease-ips/
[30] Prefixx, "White-Glove Service," 2025. [Online]. Available: https://prefixx.net/white-glove-service
[31] IPTrading, "IPv4 Address Blacklisting: Risks and Solutions," Oct. 2023. [Online]. Available: https://iptrading.com/blog/ipv4-address-blacklisting-risks-and-solutions/
[32] IPv4 Global, "Networking Education: IP Blocklist & Removal," Jun. 2020. [Online]. Available: https://ipv4.global/blog/ip-blocklist-blacklist/
[33] G. Davidavicius, "The Spamhaus Project: Step by Step Guide to Remove Your IP from Blocklists," Oct. 2024. [Online]. Available: https://www.ipxo.com/blog/remove-ip-from-spamhaus-blocklist/
[34] IPv4 Market Group, "IPv4 Blacklist Removal," 2025. [Online]. Available: https://ipv4marketgroup.com/broker-services/ipv4-blacklist-removal/
[35] IPv4 Superhub, "Blocklist Removal," Sep. 2024. [Online]. Available: https://ipv4superhub.com/blocklist-removal/
[36] IPXO, "RPKI Management at IPXO (IP Holder)," 2025. [Online]. Available: https://www.ipxo.com/kb/technical-guides/rpki-management-at-ipxo-ip-holder/
[37] Prefixx, "Lease IPv4 addresses," 2025. [Online]. Available: https://prefixx.net/lease-ipv4-addresses
[38] IPXO, "BGP hijacking protection," 2025. [Online]. Available: https://www.ipxo.com/features/bgp-hijacking-protection/
[39] RouteViews, "University of Oregon RouteViews Project," Jan. 2025. [Online]. Available: https://www.routeviews.org/routeviews/
[40] RIPE Network Coordination Centre, "Routing Information Service (RIS)," 2025. [Online]. Available: https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/
[41] Center for Applied Internet Data Analysis (CAIDA), "Inferred AS to Organization Mapping Dataset," 2025. [Online]. Available: https://www.caida.org/catalog/datasets/as-organizations/
[42] ——, "AS Relationships," 2025. [Online]. Available: https://www.caida.org/catalog/datasets/as-relationships/
[43] RIPE Network Coordination Centre, "Brokers," Dec. 2023. [Online]. Available: https://web.archive.org/web/20231216062402/https://www.ripe.net/manage-ips-and-asns/resource-transfers-and-mergers/brokers
[44] APNIC, "Registered IPv4 brokers," 2025. [Online]. Available: https://www.apnic.net/manage-ip/manage-resources/transfer-resources/transfer-of-unused-ip-and-as-numbers/transfer-facilitators/
[45] C. Tsaousis, "FireHOL IP Lists," 2025. [Online]. Available: http://iplists.firehol.org/
[46] IPv4.Global, "IP Address Leasing Hub: Monetize Your IP Addresses," 2025. [Online]. Available: https://auctions.ipv4.global/leasing/en/leasing-hub
[47] IPXO, "Market stats," 2025. [Online]. Available: https://www.ipxo.com/market-stats/
[48] B. Du, R. Fontugne, C. Testart, A. C. Snoeren, and k. claffy, "Sublet Your Subnet: Inferring IP Leasing in the Wild," in *Proceedings of the 2024 ACM on Internet Measurement Conference*, ser. IMC '24. New York, NY, USA: Association for Computing Machinery, Nov. 2024, pp. 328–336. [Online]. Available: https://dl.acm.org/doi/10.1145/3646547.3689010
[49] L. Oliver, G. Akiwate, M. Luckie, B. Du, and k. claffy, "Stop, DROP, and ROA: Effectiveness of defenses through the lens of DROP," in *Proceedings of the 22nd ACM Internet Measurement Conference*, ser. IMC '22. New York, NY, USA: Association for Computing Machinery, Oct. 2022, pp. 730–737. [Online]. Available: https://doi.org/10.1145/3517745.3561454
[50] N. Chachra, D. McCoy, S. Savage, and G. M. Voelker, "Empirically characterizing domain abuse and the revenue impact of blacklisting," in *Proceedings of the Workshop on the Economics of Information Security (WEIS)*, vol. 4, 2014. [Online]. Available: http://cseweb.ucsd.edu/~voelker/pubs/namevalue-weis14.pdf
[51] M. L. Mueller and B. Kuerbis, "Buying numbers: An empirical analysis of the IPv4 number market," in *iConference 2013*. iSchools, Feb. 2013. [Online]. Available: https://hdl.handle.net/2142/36045
[52] I. Livadariu, A. Elmokashfi, A. Dhamdhere, and k. claffy, "A first look at IPv4 transfer markets," in *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '13. New York, NY, USA: Association for Computing Machinery, Dec. 2013, pp. 7–12. [Online]. Available: https://dl.acm.org/doi/10.1145/2535372.2535416
[53] L. Prehn, F. Lichtblau, and A. Feldmann, "When wells run dry: The 2020 IPv4 address market," in *Proceedings of the 16th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '20. New York, NY, USA: Association for Computing Machinery, Nov. 2020, pp. 46–54. [Online]. Available: https://doi.org/10.1145/3386367.3431301
[54] E. Khan, E. Chiapponi, M. Verkleij, A. Sperotto, R. Van Rijswijk-Deij, and J. Van Der Ham-De Vos, "A First Look at User-Installed Residential Proxies From a Network Operator's Perspective," in *2024 20th International Conference on Network and Service Management (CNSM)*, Oct. 2024, pp. 1–9. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10814519
[55] A. Srebaliute, "Top 10 Benefits of Leasing IP Addresses for IP Holders," Oct. 2023. [Online]. Available: https://www.ipxo.com/blog/top-10-benefits-of-leasing-for-ip-holders/
[56] LARUS, "5 Strategies to Combat IP Address Abuse in the Leasing Market," Oct. 2023. [Online]. Available: https://larus.net/blog/5-strategies-to-combat-ip-address-abuse/
[57] Bailey, Michael, D. Dittrich, E. Kenneally, and Maughan, Douglas, "The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research," U.S. Department of Homeland Security, Tech. Rep., 2012. [Online]. Available: https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCORE-20120803_1.pdf

APPENDIX A
ETHICAL CONSIDERATIONS

This study follows the principles outlined in the Menlo Report [57], which extends traditional research ethics to information and communications technology (ICT) research. In line with its core principles — respect for persons, beneficence, justice, and respect for law and public interest — we carefully considered the ethical implications of our data collection and analysis. As our study did not involve human subjects, institutional review board (IRB) approval was not required.

In our interaction with representatives from IP brokers, we prioritized transparency by using our real names and affiliations. For Broker A, we used a publicly documented API authorized to an account granted to us for research purposes. We confirmed with their team that our method of access was acceptable. For Broker B, we collected data through their website. To the best of our understanding, this did not violate their terms of service. We do not disclose the names of the brokers studied to mitigate any reputational harm.

When collecting the marketplace data, we took steps to minimize any risk to the production systems of the IP brokers. To reduce system strain, we imposed a strict limit of one data fetch per day. If a request failed for any reason, we retried up to three times. If all attempts failed, we assumed the service

was temporarily unavailable and paused for an hour before making another attempt. Since our requests were modest in size and similar to routine website traffic, we assessed their impact on production systems to be negligible.

We avoided collecting and processing personally identifiable information (PII) in our study. One broker included IP holder names in their marketplace listings. Although we observed only company names, we acknowledge that personal names could appear. Nonetheless, all analysis was conducted on aggregate data only. For the dataset from Broker A, we retained only essential fields (prefix, RIR, and price) and anonymized prefixes per the broker's request. Broker B declined our request to publish data. This approach ensures that our study contributes to transparency in the IP leasing market, while upholding respect for persons and the companies' interests.

Prior to publication, we contacted both brokers to share a pre-publication version of our study and invited them to raise any concerns or offer clarifications. We carefully assessed all feedback and incorporated revisions where comments were well-founded or helped improve presentation.