A first look at the misuse and abuse of the IPv4 Transfer Market

Abstract. The depletion of the unallocated address space in combination with the slow pace of IPv6 deployment have given rise to the IPv4 transfer market, namely the trading of allocated IPv4 prefixes between ASes. While RIRs have established detailed transfer policies in an effort to regulate the IPv4 transfer market for malicious networks such as spammers and bulletproof ASes, IPv4 transfers pose an opportunity to bypass reputational penalties of abusive behavior since they can obtain "clean" IPv4 address space or offload blacklisted address space. Additionally, IP transfers create a window of uncertainty about legitimate ownership of prefixes, which leads to inconsistencies in WHOIS records and ROA objects. In this paper we provide the first detailed study of how transferred IPv4 prefixes are misused in the wild by synthesizing an array of longitudinal IP blacklists, honeypot data, and AS reputation lists compiled through hijack detection. Our findings yield evidence that the transferred network blocks are used by malicious networks to address botnets and fraudulent sites in much higher rates compared to non-transferred addresses, while the timing of the attacks indicate efforts to evade filtering mechanisms.

1 Introduction

The depletion of the unallocated IPv4 addresses combined with the slow transition to IPv6 has led to the emergence of a secondary market for ownership transfers of IPv4 addresses. However, the IPv4 market has been poorly regulated due to the lack of widely adopted IP prefix ownership authentication mechanisms, inconsistent contractual requirements between legacy and allocated address space [43], and policy incongruences among Regional Internet Registries (RIRs). As a result, IPv4 transfers have become target of fraud and abuse by malefactors who try to bypass legal IP ownership processes [19]. RIRs have responded to the emergence of the IPv4 market by establishing policy frameworks that aim to safeguard the hygiene of the accuracy of registered IP blocks and provide oversight and transparency on how organizations trade IPv4 address blocks [37, 19]. However, the effectiveness of these policies in preventing abuse of the IPv4 market remains unclear. Additionally, these policies focus only ownership and utilization issues, and they do not have provisions for malicious usage of the transferred space, for instance by bulletproof hosters who seek clean address space to address botnets and fraudulent sites. Exchanges between operators in mailing lists and messaging boards show that the operational community is worried about these dangers but still face significant difficulties when purchasing or selling address space [30, 42].

In this paper we aim to shed light on the misuse and abuse of the IP transfer market. To this end, we combine a large collection of longitudinal control-plane and data-plane data to analyze and verify the information reported by RIRs on IP transfers. We find that the reported transfer dates and recipient organizations do not reflect the state of WHOIS registries and BGP routing for more than 65% of the transferred prefixes. Additionally, 6% of the prefixes covered with ROAs have an inconsistent origin ASN. We then compile and analyze a large-scale dataset of malicious activities that covers a period of more than a decade, derived from IP traffic traces and control-plane paths, including IP blacklists, honeypots, prefix hijacking detection and AS reputation mechanisms. Our findings reveal that the transferred IP space is between 4x to 25x more likely to be blacklisted depending on the type of malicious activity, while the majority of the transferred IPs are blacklisted after the transfer date, even when the transferred address space was deployed and visible to IP scans at least a month before the transfer. The disproportionate representation of transferred prefixes in blacklists persists even when we filter-out the address space used by well-known legitimate networks, such as cloud platforms (*e.g.* Amazon Web Services, Google Cloud, Microsoft Azure) whose Infrastructure-as-a-Serive (IaaS) is often abused to host malware in short-lived Virtual Machines (VMs). Finally, we provide evidence that ASes detected to be serial BGP hijackers or bulletproof hosters are over-represented in the IPv4 market and exhibit suspicious patterns of transactions both as buyers and sellers. These results offer new insights on agile blacklist evasion techniques that can inform the development of more timely and accurate blacklisting mechanisms. Additionally, our work can inform debates on developing and evaluating RIR policies on IP transfers to improve the hygiene of the ecosystem.

2 Background and Related Work

2.1 IP Transfer Market

Today, the available IPv4 address space of all Regional Internet Registries (RIRs) except AFRINIC has been depleted [21]. Despite increasing pressure on network operators to enable IPv6, less than 30% of the ASes are currently originating IPv6 prefixes [22]. Since RIRs are unable to allocate additional IPv4 addresses, many network operators try to prolong the lifespan of IPv4 by buying address space allocated to other networks, which has led to the emergence of a secondary IP market. This market has been characterized as murky [43], due to the lack of transparency and mechanisms to authenticate the ownership of IP space.

In an effort to prevent abuse of this secondary IP market, all RIRs have devised intra-registry transfer policies, starting with RIPE in 2008 [50, 5, 8, 27, 50, 3]. All RIRs, except RIPE, have imposed restrictions on IP transactions, that require a minimum size of transferred address space and adequate justification of need from the side of buyers. Inter-regional transfers have been approved by ARIN, APNIC and RIPE. Organizations involved in such transactions have to comply with the inter-RIR transfer policies of their local registry [9, 51]. ARIN and APNIC follow the same need-based policy for intra-RIR and inter-RIR transactions. RIPE, in contrast to its intra-RIR policy, requires inter-RIR buyers to document the utilization of at least 50% of the transferred address space for five years. However, these regulations do not apply in the case of transfers that occur due to mergers and acquisitions. Under these policies, the first intra-RIR and inter-RIR transfers occurred October 2009 and October 2012, respectively. The IPv4 transfer market size has significantly increased over the years, from 17,408 to 40,463,872 IPs for intra-RIR transfers, and from 1,792 to 1,885,440 IPs for inter-RIR transfers, with the highest activity occurring within RIPE and ARIN. Moreover, 96% of the IPv4 addresses are exchanged within the same registry and most of these IP transactions occur within the North America region, while 85% of the inter-RIR transfers originate from ARIN. Despite the increasing prominence of the IPv4 market, there are only a few studies of its ecosystem. Periodically, RIRs and IPv4 address brokers report on the trends and evolution of the IP transactions [7, 46–49, 56, 1], but also a portion of buyers have reported their experiences [35, 19, 6]. Early academic studies [17, 28] discussed the possible implications of market-based mechanisms for reallocating IPv4 addresses. Mueller *et al.* [32, 33] used the list of published transfers to analyze the emerging IPv4 transfer market by quantifying the amount of legacy allocations exchanged on the markets and the impact of the need-based policies on the utilization of the transferred blocks. Livadariu *et al.* [23] provided a comprehensive study on the IPv4 transfer market evolution, the exchanged IPv4 blocks, and the impact on the routing table and IPv6 adoption. The authors also proposed a method for inferring IP transfers from publicly available data, i.e., routing advertisements, DNS names, RIR allocation and assignment data. To the best of our knowledge, no prior work has studied the IPv4 transfer market from the perspective of fraudulent behavior and misuse.

2.2 Malicious Internet Activities

An *IP blacklist* is an access control mechanism that aims to block traffic from IP addresses which have been detected to originate fraudulent activities, such as spamming, denial of service, malware or phishing. Such blacklists are compiled using spamming sinkholes, honeypots, and logs from firewalls, Intrusion Detection Systems (IDS), and anti-virus tools distributed across the Internet. Several works have studied malicious Internet activities based on IP blacklists [4, 62]. Ramachandran *et al.* [38] provided one of the first studies, by analyzing over 10 million messages received by a spam sinkhole over a period of 18 months, and by correlating them with lookups to 8 blacklists. Their results showed that combined use of blacklists detects 80% of the spamming hosts. Moreover, the network behavior of serial spammers has distinctive characteristics which can be exploited to develop behavioral and predictive blacklisting [39]. Shina *et al.* [54] evaluated the accuracy of four spamming blacklists, and found that blacklists have a very low False-Positive Rate, with 2 of the blacklists having less than 1% FPR, but high false-negative rate (above 36%) when used individually. A similar study by Kührer *et al.* [26] evaluated the effectiveness of 18 blacklists, 15 public

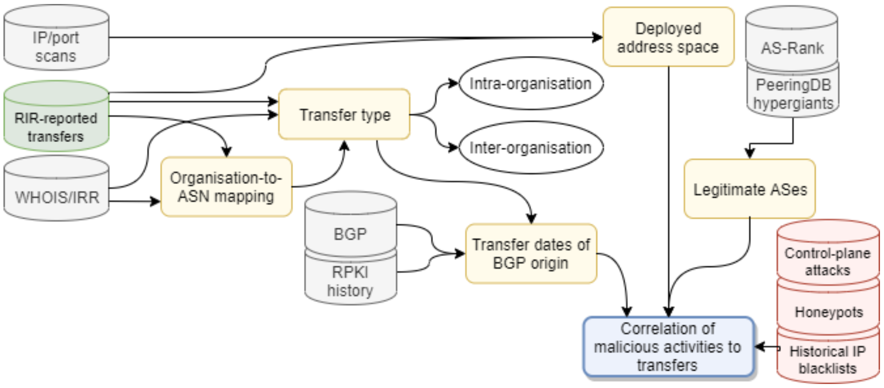

Fig. 1: Overview of datasets and measurement methodology.

and 3 by AntiVirus (AV) vendors. The authors found that blacklists derived from AntiVirus vendors are able to detect at least 70% of the malicious domains for 13 types of malware, and at least 90% for 7 of these types, outperforming significantly public blacklists. Zhao *et al.* [61] compiled an extensive historical blacklist dataset to analyze the trends and evolution of malicious activity over the span of a decade.

In addition to detecting malicious activity through monitoring data-plane traffic, an array of studies developed techniques to detect attacks through control-plane data. Shue *et al.* [53] studied the connectivity of ASes that are over-represented in 10 popular blacklists. They found that a small number of ASes with a disproportionate fraction of blacklisted address space are more likely to have dense peering connectivity and exhibit higher frequency of peering changes. Konte *et al.* proposed ASwatch [25], a system a system that aims to identify bulletproof hosting ASes by inferring irregularities in the routing and connectivity patterns. Testart *et al.* [58] profiled serial prefix hijackers by developing a supervised machine learning model, based on which they analyzed 5 years of BGP data to infer 934 ASes that exhibit persistent misbehavior.

We utilize insights and data from the above works to conduct a comprehensive analysis of malicious activity involving transferred IP prefixes and organizations that participate in the transfer market. We combine both data-plane and control-plane data to compile an extensive dataset of attacks. Our blacklist dataset combines a large number of blacklists compiled by AV vendors, which was found to be the best approach to maximize coverage and minimize false-positives.

3 Datasets and Methodology

In this section we present the data and methods that we employ to analyze malicious activities and misuse of transferred IP address space. Figure 1 shows how we synthesize and process an array of chosen datasets.

3.1 Processing of IPv4 Transfers

Collection of reported IP transfers. As a first step of our methodology we collect the list of reported intra-RIR and inter-RIR IP transfers, as published by the RIRs. For each transferred resource, we extract the IPv4 address block, the transfer date, and the names of the seller and buyer organizations. Note that none of the RIRs provides the AS Number (ASN) of the organizations involved in the transfer. In the case of inter-RIR IP transactions we also retrieve the RIR for both the seller and the buyer organizations. For intra-RIR IP transfers, ARIN and RIPE also indicate the transfer type, namely if a transfer occurred due to changes within an organization (merger and acquisitions), or as a sale of address space between distinct organizations. However, this information is not available for inter-RIR transfers and for transfers within the APNIC and LACNIC regions. Overall, we collected $30,335$ transfers involving $28,974$ prefixes between 2009-10-12 and 2019-08-24. Of these transfers, $9,564$ (31.5%) are labeled as Mergers/Acquisitions, $17,934$ (59.1%) as IP sales, while the rest $2,837$ (9.4%) are not labeled.

Mapping of organization names to ASNs. We aim next to find the ASes that map to the organizations active on the IP transfer markets. To this end we take the following steps. First, we collect historical WHOIS records every 4 months throughout the IP transfers collection period, i.e., from October 2009 to August 2019. Second, for each allocated ASN we extract the AS name and corresponding organization name for for each allocated ASN. Third, we match the organization names in the RIR transfer lists against the extracted WHOIS fields and select the corresponding ASes. We were able to map $8,744$ out of the $15,666$ organizations involved in the transfer market. Through manual inspection we found that unresolved organizations do not operate an Autonomous Systems.

Inference of transfer types. For the 9.4% of the transfers without reported transfer type (*merger/acquisition*, or *IP prefix sale*), we try to infer if the transfer occurred between siblings. For organizations we successfully mapped to ASNs, we use CAIDA's AS-to-Organization inference [13] closest to the date of the transfer. Additionally, for inter-RIR transfers and for organizations not mapped to an ASN, we compare the organization names using the string comparison algorithm introduced by Myers [34].

Correlation of transfers to BGP activity. The algorithm returns a value between 0 and 1, where 1 indicates identical strings. For values above 0.8 we consider the organizations as siblings. To improve the accuracy of string comparison, we filter-out from the organization names stop words, and the 100 most common words across all names (*e.g.* Ltd, corporation, limited). Based on

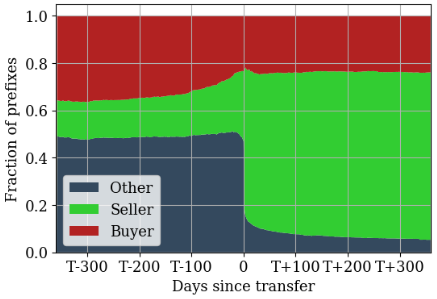

Fig. 2: Shift of origin AS in relation to transfer date

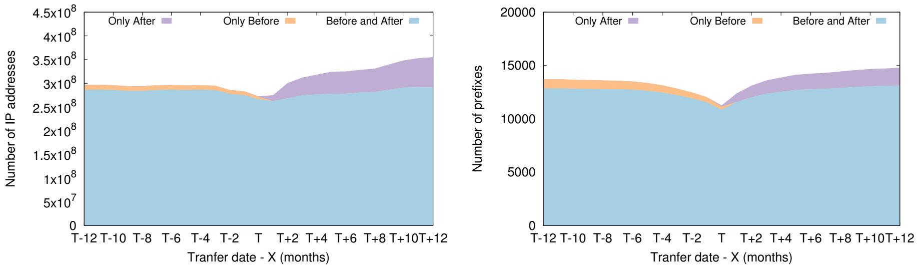

(a) Transferred IPs

(b) Transferred prefixes

Fig. 3: The visibility of transferred address space in BGP advertisments.

the above process we infer 841 (29.6%) of the unlabeled transfers to be between sibling ASes.

We use daily routing tables from all the Routeviews [16] and RIPE RIS [45] collectors, to investigate how the transferred IP address space is advertised across time. For each transfer, we check whether the transferred IP blocks are routed within one year before and one year after their reported transferred date. As shown in figure 3, 97.05% of the IPs and 64% of the prefixes are advertised consistently across the entire period. Only ≈ 10% of the prefixes are advertised only after the transfer, while about ≈ 5% are advertised only before the transfer. However, the reported transfer date does not correlate with a change in prefix origin for 65% of the transferred prefixes, while in 15% of the cases the buyer advertises the prefix one year before the transfer.

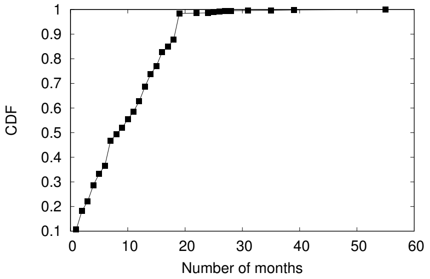

Fig. 4: Distribution of the origin inconsistencies for ROA records covered by transferred address space.

Inconsistencies in the routing advertisements. We also investigate the existence of RPKI invalid advertisement. To this end, we run ziggy [36] to collect RPKI data for January and July of each year between 2015 to 2019. We then search in the collected data for prefixes that match the transferred address space. For such prefixes we further compare for each month the origin AS with the registered ASN in the collected RPKI object. We label as *origin inconsistencies* case where we discover a mismatch between the two ASes. We collect an overall of 15,663 ROA records that are covered by 5476 IP transferred prefixes. However, we find origin inconsistencies for only 951 of the records. Figure 4 shows the distribution of the number of months for which we detect such inconsistencies. Surprisingly, we find that for 40% of the prefixes such advertisements last more than one year, while some cases last more than three years.

Measuring the deployed transferred IP space. The behavior of BGP paths will help us interpret more accurately the observed malicious activities, nonetheless routed address space is not necessarily deployed and used in practice [15, 44]. To study the malicious behavior of the deployed transferred address space we collect Internet-wide IP scans every 3 months between 2012-01-02 and 2019-09-01. We first collect the `ICMP ECHO REQUEST` scans from the USC/ISC project LANDER [18], which sweeps the IANA allocated IP ranges, and records all the IPs that respond with an `ICMP ECHO REPLY` message. We complement these data with Internet-wide UDP and TCP scans collected by RAPID7's project Sonar [41, 40], which records the IPs that respond to ZMAP probes against popular UDP and TCP services.

3.2 Detection of malicious IPs and ASes

After we compile and process the IP transfers, we construct an extensive dataset of cyber-attack sources to analyze the hygiene of the transferred address blocks and the players within the IPv4 market.

Real-time BlackLists (RBLs) provide one of the most popular techniques to detect networks responsible for attacks. Unfortunately, most blacklist providers do not offer historical snapshots, but typically they only publish the blacklist at a certain web location that is refreshed periodically – daily or even hourly – so that firewalls can automatically update their rules. However, we were able to find two large-scale historical blacklist datasets compiled and archived by third-parties.

FinalBlacklist. Zhao *et al.* [61] compiled the `FinalBlacklist` dataset that contains over 51 million blacklisting reports for 662K IPs between January 2007 and June 2017, as part of a decade-long analysis of malicious Internet activity. To construct the `FinalBlacklist`, the authors collected historical blacklist snapshots through the Wayback Machine [24], which they extended using VirusTotal [60], an API that aggregates periodic data from more than 70 Anti-Virus and blacklist data feeds. 7.6 million (15%) of the blacklisting reports is labeled by the original source with the type of the malicious activity, which the authors abstract into six classes: Exploits, Malware, Fraudulent Services (FS), Spammers, Phishing, and Potentially Unwanted Programs (PUP). Based on the labeled subset they employed a random forest classifier to predict the class of the remaining 44M blacklisted activities with 92.4% accuracy. 90.9% of the blacklisted IPs correspond to malware, while only (0.01%) correspond to Spammers.

RIPE Stat Abuse API. To augment the `FinalBlacklist` dataset with IPs involved in the distribution of Spamming, we rely on data published by RIPE NCC who is archiving daily snapshots since 2009-06-22 of the `UCE-Protect` Network [59] blacklist[1], one of the most prominent anti-spamming blacklists. RIPE NCC provides public access to these data through the RIPE Stat RESTful API [52], which allows querying the blacklisting reports for a specific IP prefix (no bulk querying). If an IP range within the queried prefix is blacklisted, the API returns the blacklisting period (start and end date), allowing us to collect historical blacklisting reports.

[1] RIPE Stat also provides access to Spamhaus DROP snapshots [55], which we do not use because it covers only directly allocated address space

The `UCE-Protect` blacklist uses three different levels of blacklisting policies, according to the severity and persistence of the observed malicious activity. *Level-1* blacklists only single IP addresses detected to deliver e-mails to spam traps, conduct port scans or attack the `UCE-Protect` servers. Level-1 records expire automatically after 7 days if there are no further attacks. *Level-2* aims to stop escalating attacks by blacklisting IP prefixes with multiple IPs that emanate spam repeatedly for a week, implying lack of appropriate security measures or intentional misbehaviour. *Level-3* blacklists all IPs within an ASN if more than 100 IPs, but also a minimum of 0.2% of all IPs allocated to this ASN, are Level-1 blacklisted within 7 days. This aggressive policy assumes that legitimate networks are unlikely to have a sustained high volume of blacklisted IPs. Additionally, a prefix/ASN can get Level-2/3 blacklisted if a network employs evasion techniques against blacklists, such as rotating the IPs of abusers within a prefix, or blocking IP addresses of blacklist providers.

Detection of persistent C&C hosters. The activity of botnets is typically coordinated by Command and Control (C&C) servers. C&C servers may only orchestrate and not participate in attacks themselves, therefore their detection is primarily based on honeypots. Shutting down of C&C servers is critical in defending against botnets, an effort that may even involve security agencies such as the FBI [31], therefore legitimate network operators tend to respond quickly in requests for C&C take-downs in contrast to bulletproof hosters. We use data from two distributed honeypots operated by BadPackets [10] and BinaryEdge [11] to detect ASes that host C&C servers for over two weeks, despite notices by the honeypot operators. We were able to detect 28 ASes that are persistent and serial C&C hosters between February 2018 and June 2019.

AS reputation lists based on BGP misbehavior We complement the set of malicious ASes compiled through the honeypot data with AS reputation lists which are developed by monitoring the BGP routing system to detect ASes with consistent patterns of malicious routing, such as traffic misdirection. We use the list produced by Testart *et al* [58], which we further extend with examples of bulletproof hosters and hijackers reported by [25, 14] resulting in a list of 922 malicious ASes.

4 Analysis and Results

Blacklisted Address Space. We first compare the malicious activity emanating from transferred and non-transferred prefixes as reflected by our IP blacklist reports. Table 1 summarizes the blacklist records per type of malicious activity, for transferred and non-transferred IPs and prefixes. Transferred IPs are disproportionately represented in the blacklist for every type of malicious activity except Spamming. In particular, the transferred address space represents only 16% of the total address space[2], but covers 61% of the blacklisted IPs. The fraction of transferred prefixes with at least one blacklisted IP is 4x to 25x larger

[2] Same percentage when we take into account only routed address space

Table 1: Analysis of blacklisted IPs. Transferred IP prefixes are disproportionately represented in all the blacklists by a rate between 4x for Malware IPs, to 43x for Fraudulent services.

Blacklist type	Blacklisted IPs part of transfers		Trans Prefixes w/ blacklisted IPs		Non-trans Prefixes w/ blacklisted IPs
	All	Filtered	All	Routed	Routed
Unwanted Programs	55%	43%	3.6%	5.5%	0.95%
Exploits	30%	30%	4.7%	7.2%	0.92%
Malware	36%	29%	16.6%	25.3%	6.2%
Phising	36%	25%	7.5%	11.6%	2%
Fraudulent Services	23%	27%	3.8%	9.6%	0.22%
Spammers	12%	12%	0.6%	0.9%	0.1%

than the fraction of non-transferred prefixes for every blacklist type, with Spam being the category with the smallest fraction of blacklisting reports per prefix.

As shown in Figure 5, 40% of the routed transferred prefixes appear at least once in our RBLs, compared to only 6% of the non-transferred routed prefixes. However, the blacklisting activity does not originate uniformly across the address space. When we break down all prefixes to their covered /24 sub-prefixes we find that the blacklisted IPs are concentrated in 6% of the transferred /24s, and in only 1% of the non-transferred /24s (Figure 5b). This happens because some of the less specific transferred prefixes are owned by large-scale legitimate networks, such as Tier-1 providers, that proportionally originate a very small fraction of blacklisting reports. For example, the prefix `50.128.0.0/9` which was transferred by an acquisition from Comcast includes 32, 768 /24 sub-prefixes (more than all transferred prefixes), but has only 289 blacklisting reports. Still, transferred /24 sub-prefixes are 6x more likely to be blacklisted, than the non-transferred ones.

Blacklisted ASNs. We analyze the blacklisting reports per ASN, to understand how the detected malicious activity is distributed across the participants of the IP transfer market. Almost 50% of all the ASNs that participate in the transfer market appear at least once in the blacklist, compared to only 16% of the ASNs that do not participate in the transfer market and appear in the BGP table to originate prefixes (Figure 6a). Moreover, ASes in the transfer market tend to have larger fraction of their address space blacklisted, with a median of 0.06% compared to 0.03% for ASes not involved in any transfer, which is an indication of more consistent malicious behaviour. This trend is even more pronounced for ASes that are both sellers and buyers of IP prefixes, which for some ASes appear to be a strategy to recycle blacklisted prefix. To study whether the

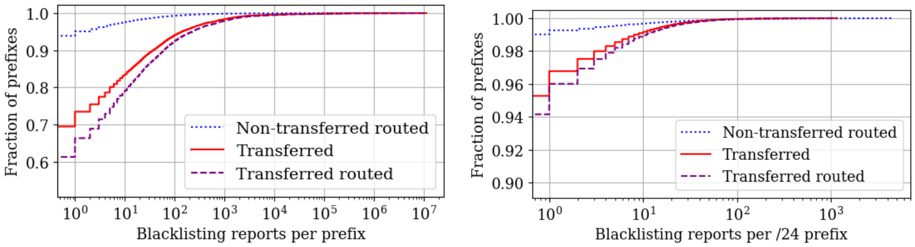

(a) Blacklist records per prefix

(b) Blacklist records per /24

Fig. 5: Distribution of the volume of blacklisting reports for transferred and non-transferred prefixes.

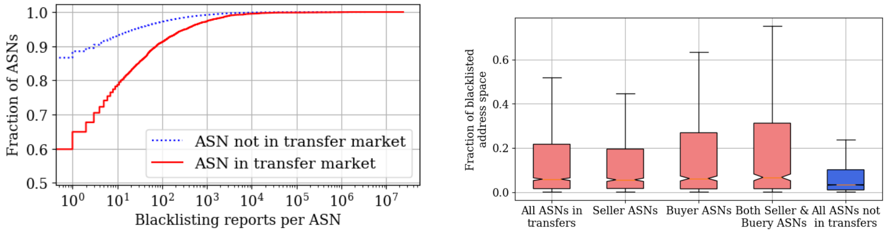

(a) Distribution of blacklisting reports per ASNs (b) Fraction of blacklisted address space per ASN

Fig. 6: Comparison of the blacklisted activity of the ASNs in the transfer market, compared to the rest of the ASNs that originate BGP-routed prefixes.

higher blacklisting rate of ASNs involved in transfers may be explained by a bias in the composition of ASNs that exchange IP space, we compare their user population according to APNIC's estimates [20], and we also compare their self-reported business type in PeeringDB [2]. For both datasets the composition of ASNs is very similar, with ASNs absent from transfers exhibiting slightly higher median user population.

While the blacklisted prefixes are distributed across half of the ASNs involved in transfers, there are 26 ASes with more than 10K blacklisted IPs, including prominent cloud providers (*e.g.* Amazon, Microsoft, Google, OVH) and Tier-1 providers (*e.g.* GTT, CenturyLink, Seabone). Attackers often utilize cloud platforms as a cost-effective way to host malicious software in short-lived Virtual Machines and avoid detection [57], while large providers operate a global network that covers a massive user population. These ASes account for only 0.3% of the blacklisted prefixes, but cover 55% of all the blacklisted IPs, which explains the the long tail of the distributions in Figure 5a. Following a filtering approach similar to the one proposed by Testart *et al.* [58], we consider as non-suspicious the 1,000 ASes with the largest customer cones according to AS-Rank [29]. However, cloud providers, CDNs and large-scale eyeballs have relatively small customer cones. Therefore, we complement the filtered ASes with: *(i)* the 30 ASes with the largest amount of traffic (*hypergiants*) based on the methodology by Böttger *et al.* [12], and the 1000 ASes with the largest user population according to APNIC. As shown in the column *"Filtered"* of Table 1, even when filtering out these ASes, the fraction of blacklisted transferred IPs is between 2x − 3x higher than the total fraction of transferred IPs, while the fraction of blacklisted prefixes is virtually identical between the filtered and the non-filtered datasets. This is an indication that *a large number of ASes in the transfer market exhibit higher affinity for malicious activity which cannot be explained by their business model network footprint.* The rest of our analysis is based on the filtered set.

Blacklisting timing. To explore the dynamics between malicious activity and the IP transfers, we compare the timing of the blacklisting reports to the transfer date. We use the effective transfer date, as observed by BGP routing changes (see Section 3.1), and the reported transfer time only when the origin AS does not change at all. As shown in figure 7, the number of blacklisted IPs peaks within a year of the transfer date for all types of malicious activity. Such blacklist-

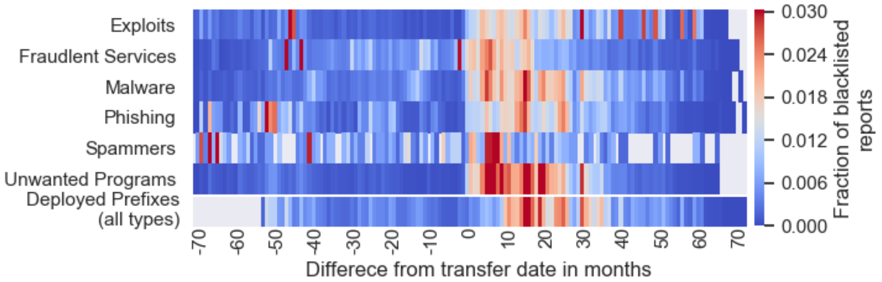

Fig. 7: Blacklist reports per type of malicious activity for transferred IPs, compared to the transfer date. The last row shows the blacklisting activity for deployed prefixes based on the Internet-wide IP and port scans.

ing activity shortly after the transfer date may happen because the transferred addresses were unused before the transfer.

To illuminate this possibility, in the last row of figure 7 we plot the blacklisting reports only for prefixes with IPs visible in our IP/port scans at least one month before the transfer date. For deployed prefixes the peak in malicious activity also peaks after the transfer date, but after one year. This finding indicates that recipients of IP addresses are more prone to abuse of the IP space, which agrees with the difference in blacklisting magnitude between buyers and senders as shown in figure 6b.

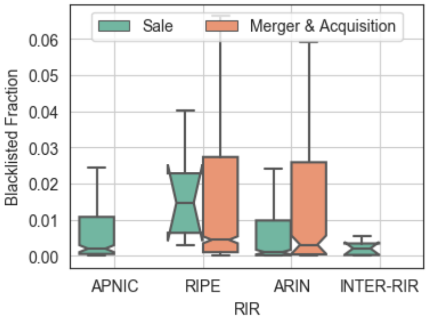

Fig. 8: Fraction of address space in IP blacklists per transfer region, and per type of transfer

Per-region and per-transfer type differences We then investigate whether the malicious activity differs between regions and between transfer types. Figure 8 compares the fraction of blacklisted transferred address space between prefixes exchanged as Merge & Acquisitions and as IP sales for each region with blacklisted IPs, and for inter-region transfers. Prefixes exchanged within the RIPE region as sales originate have the highest fraction of blacklisted IPs, which is statistically significant.

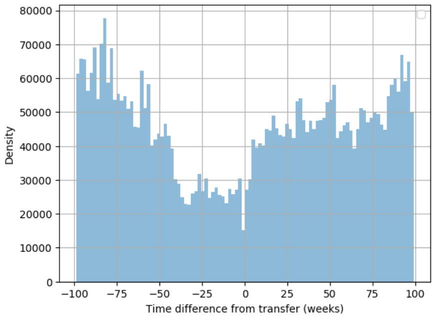

Fig. 9: Density of blacklisted IPs for low-reputation ASes that participate both as buyers and sellers in the IPv4 market

In contrast, ARIN exhibits higher malicious activity from prefixes transferred between siblings, although the spread of values makes it difficult to generalize this observation. For APNIC and inter-RIR transfers we observe only non-sibling blacklisted transactions, while for AFRINIC and LACNIC we do not have any blacklisted transferred IP (after the AS filtering step).

Participation of low-reputation ASes in IPv4 transfers The final part of our analysis is to check the participation rate of low-reputation ASes (hijackers, C&C and bulletproof hosters) in IP transfers. Although 85% of the ASes visible in the BGP routing table are not involved in IP transfers, 47% of the low-reputation ASes have been either buyers (48%) or sellers (52%). Surprisingly, 32% of these ASes participate both as buyers and sellers. This practice may signal an attempt to recycle "tainted" address space in order to evade blacklist filters, since blacklist providers may remove listed IPs and prefixes when there is a shift in ownership. Figure 9 shows that indeed the density of blacklisted IPs for the low-reputation buyer/seller ASes dips at the transfer date and increases shortly thereafter.

5 Conclusion

In this paper we present a first comprehensive measurement study of malicious activities within the transferred IPv4 address space and the networks that are involved in the IPv4 market. We first combine a wide range of control-plane and data-plane data to process the details of the reported IP transfer reports and and verify the ownership of the exchanged prefixes based on BGP paths and historical WHOIS and RPKI data. We find that for more than 65% of the IP transfers, the origin ASes and the transaction dates appear to be inconsistent with the transfer reports, while 6% of ROAs become stale after the transfer for many months. Our results reveal at best poor practices of resource management that can facilitate malicious activities, such as hijacking attacks, and even lead to connectivity issues due to the increasing deployment of RPKI-based or IRR-based filtering mechanisms.

We then analyze the exchanged IPv4 address blocks against an extensive dataset of malicious activities that span a decade, which includes IP blacklists, honeypot data, and nonlegitimate ASes based on the detection of control-plane misbehavior. Our findings show that the ASes involved in the transfer market exhibit consistently higher malicious behavior compared to the rest of the ASes, even when we account for factors such as business models and network span. Our findings are likely to be a lower bound of malicious activity from within transferred IP addresses since a number of transactions may occur without being reported to the RIRs. As part of our future work we aim to extend our analysis to non-reported IPv4 transfers and develop predictive techniques for blacklisting based on the monitoring of the IPv4 transfer market.

We believe that these insights can inform the debates and development of RIR policies regarding the regulation of IPv4 markets, and help operators and brokers conduct better-informed due diligence to avoid misuse of the transferred address space or unintentionally support malicious actors. Moreover, our results can provide valuable input to blacklist providers, security professionals and researchers who can improve their cyber-threat monitoring and detection approaches, and tackle evasion techniques that exploit IPv4 transfers.

References

1. IPv4 Market Group, https://ipv4marketgroup.com/broker-services/buy/
2. PeeringDB. https://www.peeringdb.com
3. AFRINIC: IPv4 Resources transfer within the AFRINIC Region, https://www.afrinic.net/policy/manual#IPv4-Resource-Transfers-within-AFRINIC-Region
4. Alieyan, K., ALmomani, A., Manasrah, A., Kadhum, M.M.: A survey of botnet detection based on dns. Neural Computing and Applications **28**(7), 1541–1558 (2017)
5. APNIC: APNIC transfer, merger, acquisition, and takeover policy (2010), https://www.apnic.net/policy/transfer-policy_obsolete
6. APNIC blog ,D. Huberman: Seven steps to successful IPv4 transfers (2017)
7. APNIC blog, G. Huston: IPv4 Address Exhaustion in APNIC (2015), https://blog.apnic.net/2015/08/07/ipv4-address-exhaustion-in-apnic/
8. ARIN: ARIN Number Resource Policy Manual (Version 2010.1) (2009), https://www.arin.net/policy/nrpm.html
9. ARIN: ARIN Number Resource Policy Manual (Version 2012.3) (2012), https://www.arin.net/policy/nrpm.html
10. BadPackets: Cyber-Threat Intelligence: Botnet C2 Detections. https://badpackets.net/botnet-c2-detections/ (2019)
11. BinaryEdge: HoneyPots / Sensors. https://www.binaryedge.io/data.html (2019)
12. Böttger, T., Cuadrado, F., Uhlig, S.: Looking for hypergiants in peeringdb. ACM SIGCOMM Computer Communication Review **48**(3), 13–19 (2018)
13. CAIDA: Inferred AS to Organization Mapping Dataset. http://www.caida.org/data/as_organizations.xml
14. Cho, S., Fontugne, R., Cho, K., Dainotti, A., Gill, P.: Bgp hijacking classification. In: 2019 Network Traffic Measurement and Analysis Conference (TMA). pp. 25–32. IEEE (2019)
15. Dainotti, A., Benson, K., King, A., claffy, k., Kallitsis, M., Glatz, E., Dimitropoulos, X.: Estimating internet address space usage through passive measurements. SIGCOMM Comput. Commun. Rev. **44**(1), 42–49 (Dec 2013)
16. David Meyer: University of Oregon Route Views Project, http://www.routeviews.org/
17. Edelman, B.: Running out of numbers: Scarcity of ip addresses and what to do about it. In: Auctions, Market Mechanisms and Their Applications (AMMA) (2009)
18. Heidemann, J., Pradkin, Y., Govindan, R., Papadopoulos, C., Bartlett, G., Bannister, J.: Census and survey of the visible internet. In: Proceedings of the ACM Internet Measurement Conference. pp. 169–182. ACM (Oct 2008)
19. Huberman, D.: Smarter Purchasing of IPv4 Addresses from the Market. APNIC 2017. https://2017.apricot.net/assets/files/APIC674/apricot-2017-huberman_1487269883.pdf (March 2017)
20. Huston, G.: How Big is that Network? https://labs.apnic.net/?p=526 (October 2014)
21. Huston, G.: Ipv4 address report. https://ipv4.potaroo.net/ (October 2019)
22. Huston, G.: Ipv6 / ipv4 comparative statistics. http://bgp.potaroo.net/v6/v6rpt.html (October 2019)
23. I. Livadariu and A. Elmokashfi and A. Dhamdhere: On IPv4 Transfer Markets: Analyzing reported transfers and inferring transfers in the wild. Computer Communications **111**, 105 – 119 (2017)

24. Internet Archive: Wayback Machine. https://archive.org/web/ (2001)
25. Konte, M., Perdisci, R., Feamster, N.: ASwatch: An as reputation system to expose bulletproof hosting ASes. ACM SIGCOMM Computer Communication Review **45**(4), 625–638 (2015)
26. Kührer, M., Rossow, C., Holz, T.: Paint it black: Evaluating the effectiveness of malware blacklists. In: International Workshop on Recent Advances in Intrusion Detection. pp. 1–21. Springer (2014)
27. LACNIC: One-way interregional transfers to LACNIC (2017), https://politicas.lacnic.net/politicas/detail/id/LAC-2017-2?v=2&language=EN
28. Lehr, W., Vest, T., Lear, E.: Running on empty: the challenge of managing internet addresses. In: Research Conference on Communications, Information and Internet Policy (2008)
29. Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., et al.: As relationships, customer cones, and validation. In: Proceedings of the 2013 conference on Internet measurement conference. pp. 243–256. ACM (2013)
30. Matt Torres: Purchasing IPv4 space - due diligence homework. NANOG mailing list. https://mailman.nanog.org/pipermail/nanog/2019-April/100390.html (March 2018)
31. McMillen, D.: The inside story on botnets. IBM X-Force Research (September 2016)
32. Mueller, M., Kuerbis, B.: Buying numbers: An empirical analysis of the ipv4 number market. In: Proceedings of iConference (2013)
33. Mueller, M., Kuerbis, B., Asghari, H.: Dimensioning the elephant: An empirical analysis of the ipv4 number market. In: GigaNet: Global Internet Governance Academic Network, Annual Symposium (2012)
34. Myers, E.W.: An O(ND) difference algorithm and its variations. Algorithmica **1**(1-4), 251–266 (1986)
35. NANOG 68, A. Potter: How to Navigate Getting IPv4 Space in a Post-Run-Out World (2017)
36. NLnet Labs Blog: Leaping through RPKI history with Ziggy. https://medium.com/nlnetlabs/leaping-through-rpki-history-with-ziggy-64523231a31c
37. Nobile, L.: Whois accuracy. ARIN 39. (April 2017)
38. Ramachandran, A., Feamster, N.: Understanding the network-level behavior of spammers. In: ACM SIGCOMM Computer Communication Review. vol. 36, pp. 291–302. ACM (2006)
39. Ramachandran, A., Feamster, N., Vempala, S.: Filtering spam with behavioral blacklisting. In: Proceedings of the 14th ACM conference on Computer and communications security. pp. 342–351. ACM (2007)
40. RAPID7: Project Sonar TCP Scans. RAPID7 Open Data. https://opendata.rapid7.com/sonar.tcp/ (2019)
41. RAPID7: Project Sonar UDP Scans. RAPID7 Open Data. https://opendata.rapid7.com/sonar.udp/ (2019)
42. Reddit Networking: What are your experiences with the IPv4 secondary market? https://tinyurl.com/yyumhax5 (March 2018)
43. Richter, P., Allman, M., Bush, R., Paxson, V.: A primer on ipv4 scarcity. SIGCOMM Comput. Commun. Rev. **45**(2), 21–31 (Apr 2015). https://doi.org/10.1145/2766330.2766335, http://doi.acm.org/10.1145/2766330.2766335

44. Richter, P., Smaragdakis, G., Plonka, D., Berger, A.: Beyond counting: new perspectives on the active ipv4 address space. In: Proceedings of the 2016 Internet Measurement Conference. pp. 135–149. ACM (2016)
45. RIPE: Routing Information Service (RIS), http://www.ripe.net/ris/
46. RIPE Labs, R. Wilhem: Developments in IPv4 Transfers (2016), https://labs.ripe.net/Members/wilhelm/developments-in-ipv4-transfers
47. RIPE Labs, R. Wilhem: Impact of IPv4 Transfers on Routing Table Fragmentation (2016), https://labs.ripe.net/Members/wilhelm/impact-of-ipv4-transfers-on-routing-table-fragmentation
48. RIPE Labs, R. Wilhem: Trends in RIPE NCC Service Region IPv4 Transfers (2017), https://labs.ripe.net/Members/wilhelm/trends-in-ipv4-transfers
49. RIPE Labs, R. Wilhem: A Shrinking Pie? The IPv4 Transfer Market in 2017 (2018), https://labs.ripe.net/Members/wilhelm/a-shrinking-pie-the-ipv4-transfer-market-in-2017
50. RIPE NCC: Intra-RIR Transfer Policy Proposal (2012), https://www.ripe.net/participate/policies/proposals/2012-03
51. RIPE NCC: Inter-RIR Transfers (2015), https://www.ripe.net/manage-ips-and-asns/resource-transfers-and-mergers/transfers/inter-rir-transfers
52. RIPE NCC: RIPE Stat Data API: Blacklists. https://stat.ripe.net/docs/data_api#blacklist (2019)
53. Shue, C.A., Kalafut, A.J., Gupta, M.: Abnormally malicious autonomous systems and their internet connectivity. IEEE/ACM Transactions on Networking (TON) **20**(1), 220–230 (2012)
54. Sinha, S., Bailey, M., Jahanian, F.: Shades of grey: On the effectiveness of reputation-based blacklists. In: 2008 3rd International Conference on Malicious and Unwanted Software (MALWARE). pp. 57–64. IEEE (2008)
55. Spamhaus: Don't Route Or Peer List (DROP). https://www.spamhaus.org/drop/
56. Streambank, H.: IPv4 Auctions, https://auctions.ipv4.global/
57. Technologies, W.: Internet Security Report: Q2 2019 (September 2019)
58. Testart, C., Richter, P., King, A., Dainotti, A., Clark, D.: Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. In: Proceedings of the Internet Measurement Conference. pp. 420–434. ACM (2019)
59. UCEPROTECT-Orga: UCEPROTECT Network Project. http://www.uceprotect.net/en/
60. VirusTotal: Free Online Virus Malware and URL scanner. https://www. virustotal. com/en (2012)
61. Zhao, B.Z.H., Ikram, M., Asghar, H.J., Kaafar, M.A., Chaabane, A., Thilakarathna, K.: A decade of mal-activity reporting: A retrospective analysis of internet malicious activity blacklists. In: Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security. pp. 193–205. Asia CCS '19, ACM (2019)
62. Zhauniarovich, Y., Khalil, I., Yu, T., Dacier, M.: A survey on malicious domains detection through dns data analysis. ACM Computing Surveys (CSUR) **51**(4), 67 (2018)