

Towards A User-Level Understanding of IPv6 Behavior

Frank Li

Georgia Institute of Technology / Facebook*
frankli@gatech.edu

David Freeman

Facebook
dfreeman@fb.com

ABSTRACT

IP address classification and clustering are important tools for security practitioners in understanding attacks and employing proactive defenses. Over the past decade, network providers have begun transitioning from IPv4 to the more flexible IPv6, and a third of users now access online services over IPv6. However, there is no reason to believe that the properties of IPv4 addresses used for security applications should carry over to IPv6, and to date there has not yet been a large-scale study comparing the two protocols at a user (as opposed to a client or address) level.

In this paper, we establish empirical grounding on how both ordinary users and attackers use IPv6 in practice, compared with IPv4. Using data on benign and abusive accounts at Facebook, one of the largest online platforms, we conduct user-centric analyses that assess the spatial and temporal properties of users' IP addresses, and IP-centric evaluations that characterize the user populations on IP addresses. We find that compared with IPv4, IPv6 addresses are less populated with users and shorter lived for each user. While both protocols exhibit outlying behavior, we determine that IPv6 outliers are significantly less prevalent and diverse, and more readily predicted. We also study the effects of subnetting IPv6 addresses at different prefix lengths, and find that while /56 subnets are closest in behavior to IPv4 addresses for malicious users, either the full IPv6 address or /64 subnets are most suitable for IP-based security applications, with both providing better performance tradeoffs than IPv4 addresses. Ultimately, our findings provide guidance on how security practitioners can handle IPv6 for applications such as blocklisting, rate limiting, and training machine learning models.

CCS CONCEPTS

• **Networks** → **Network layer protocols**; • **Security and privacy** → **Network security**.

ACM Reference Format:

Frank Li and David Freeman. 2020. Towards A User-Level Understanding of IPv6 Behavior. In *ACM Internet Measurement Conference (IMC '20)*, October 27–29, 2020, Virtual Event, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3419394.3423618>

*The author was a visiting researcher at Facebook at the time of this work.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
IMC '20, October 27–29, 2020, Virtual Event, USA
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8138-3/20/10.
<https://doi.org/10.1145/3419394.3423618>

1 INTRODUCTION

Attackers on the internet steal sensitive data, hijack remote machines, compromise user accounts, or perform various other harmful actions. However, all remote attacks have a common property: they must come from one or more IP addresses. Because of this property, IP address classification and clustering are important tools for security practitioners. IP address analysis helps us isolate or attribute attacks, and knowing which IP addresses are more likely to be the sources of attacks helps us improve proactive defenses.

Until recently, nearly all internet traffic used IPv4, which uses 32-bit addresses. Due to overcrowding in this small address space, over the last decade network providers have increasingly moved to IPv6, which uses 128-bit addresses. Today, a third of users now access online services using IPv6 [19]. However, with this transition, there is no reason to believe that the properties of IPv4 addresses used for security applications should translate to IPv6. Thus to guide the security operations and decisions of online platforms, it is important that we establish empirical grounding on how both benign users and attackers use IPv6 in practice.

There does exist a body of work on IPv6 use in the wild, with studies that quantify IPv6 adoption rates [25, 39], develop methods for enumerating active IPv6 addresses [15, 17, 27], explore the IPv6 internet's topology [1, 9], and assess IPv6 address allocation dynamics [25, 30]. Many of these works have studied IPv6 use by servers or the IPv6 internet at large. Work focused on IPv6 use by end users has typically characterized the IPv6 capabilities of client software or the dynamics of the client IPv6 addresses themselves. The common thread of these prior studies is that they treat IPv6 addresses or clients (and their traffic) as distinct independent units, and analyze their populations. However, online platforms often operate and make decisions at a user account granularity. For example, online platforms use machine learning models to detect malicious users, and make IP blocklisting or rate limiting decisions by weighing the tradeoff between blocking abusive users and inflicting collateral damage on benign ones. Thus there is a salient need to understand IPv6 behavior when considering users as the population units.

Our contribution. In this paper, we expand upon prior work to investigate IPv6 behavior at the user level, from both user-centric and IP-centric perspectives. From a user-centric perspective, we consider the spatial and temporal properties of the IPv6 addresses associated with a user over time, whereas prior work treated each address as independent. From an IP-centric perspective, we assess the user populations on IPv6 addresses, whereas previous studies treated addresses as atomic units, agnostic to the fact that multiple users may reside on an address. We seek to answer four primary research questions:

- **RQ1:** How does user behavior compare across IPv4 and IPv6?
- **RQ2:** How do attackers behave on each protocol?

- **RQ3:** What are the characteristics of outlier users and IP addresses?
- **RQ4:** How should we treat IPv6 addresses in security applications (e.g., blocklisting, rate limiting)?

We explore these questions using data on benign and abusive accounts at Facebook, one of the largest online platforms with more than 1.7 billion daily active users during the period studied [12]. In Sections 5 and 6 we explore RQ1–3 from the user-centric and IP-centric directions, respectively. We examine RQ4 in Section 7.

Our results. We find that in aggregate, IPv6 addresses are less populated with users than IPv4 addresses, and are much shorter lived for each user. We observe outlying behavior on both protocols, such as users who obtain thousands of IP addresses in a week, and IP addresses with hundreds of thousands of users. However, we determine that compared with IPv4 outliers, IPv6 outliers are significantly less prevalent, less diverse, and more readily predicted.

We also investigate the effects of subnetting IPv6 addresses at different prefix lengths, and we find that while /56 subnets are closest in behavior to IPv4 addresses for malicious users, either full IPv6 addresses or /64 subnets are most suitable for IPv6-based security applications, with both providing higher recall and fewer false positives than with IPv4 addresses. These insights ultimately provide guidance on how security practitioners can handle IPv6 differently than IPv4 for applications such as blocklisting, rate limiting, and training machine learning models.

Finally, we note that our study occurred during the global COVID-19 pandemic, when many people drastically shifted their physical behavior. While not a goal of our study, our dataset characterization in Section 4 revealed that this event shifted global IPv6 behavior, although only to a minor extent.

2 RELATED WORK

Despite being proposed over 20 years ago [8], IPv6’s internet-wide adoption has been a long and slow process. However, over the last 5–10 years, IPv6 deployment levels have increased significantly [19], fueling interest in understanding IPv6 use in practice. Here, we outline relevant related works on IPv6 use in both benign and adversarial settings.

IPv6 Behavior. As we are studying the IPv6 behavior of end users, we particularly focus on prior work investigating IPv6 clients. However, we note that researchers have also studied IPv6 servers [15, 17, 27], the IPv6 internet’s topology [1, 9], and IPv4 address dynamics [29, 37].

In the early days of IPv6 adoption, one body of work examined the IPv6 capabilities of clients. In 2009, Karpilovsky et al. [23] explored quantifying IPv6 deployment using BGP and ISP traffic data, observing different metrics depending on the data source, but that overall deployment remained experimental. Colitti et al. [6] measured IPv6 client adoption in 2010 from Google’s perspective, also finding that adoption remained low but was growing through a small number of large IPv6 deployments. Zander et al. [39] developed web-based techniques to measure IPv6 client capabilities in 2011 and 2012, finding that the majority of dual-stack clients still preferred IPv4 and exhibited frequent use of happy eyeballs (fast fail-over from IPv6 to IPv4). Malone [25] analyzed client addresses

to infer how clients were using IPv6, such as through IPv4-to-IPv6 transition protocols.

More recent studies investigated the structure of active IPv6 addresses, particularly to understand address allocation strategies for scanning and measurement purposes. Foremski et al. [15] used entropy analysis, clustering, and statistical modeling on collected IPv6 addresses to infer network address layouts. Similarly, Murdock et al. [27] demonstrated that active IPv6 addresses often cluster within dense regions of the address space. Gasser et al. [17] further identified addressing schemes that can be used to identify active IPv6 hosts. In a work that hews closest to ours, Plonka and Berger [30] characterized the temporal and spatial distribution of active IPv6 addresses, finding that the majority of addresses are ephemeral and that /64 prefixes are particularly dense with addresses.

These prior works provided valuable insights into the dynamics of active IPv6 addresses and clients. However, they treated IP addresses and clients as independent atomic units of analysis. As online platforms often operate and make security decisions at the user granularity, it is important to expand our understanding of IPv6 behavior to the user level. Our work takes a step in this direction.

IP-Based Security Applications. IPv4 addresses have been widely used for understanding and detecting attacks, such as with fake and compromised online accounts [2, 38], spam and other online abuses [31, 33–36], and distributed denial-of-service attacks [32]. However, until recently, IPv6 has received little attention from a security perspective, in large part due to its limited adoption. There have been some initial studies on IPv6-specific abuses, such as investigating IPv6 internet background radiation [7] and scanners [16]. Our work provides a deeper exploration into attacker IPv6 behavior, particularly for abuse on online platforms.

3 EMPIRICAL DATASET

In this section, we detail the data we use for our study as well as the limitations of our approach.

3.1 Data Collection

Our study aims to characterize IP address usage at the user level for both benign users and attackers. We do so from the perspective of Facebook, one of the largest online platforms. While the perspective of any particular platform is biased by its user base, Facebook has more than 1.7 billion daily active users around the world [12] and supports global IPv6 access (i.e., all authoritative DNS nameservers serve AAAA records), so we believe that studying IP usage from Facebook’s vantage point can provide generalizable insights.

We collect telemetry specifically for our study from external HTTP requests to Facebook servers made by logged-in users, including requests made to Facebook’s web (including the mobile-friendly version) and API (e.g., from Android and iOS apps) endpoints. For each request, we collect the following telemetry.

- The request timestamp.
- The user ID for the logged-in user making the request.
- The request’s source IP address.
- The ASN associated with the source IP address.

- The country-level geolocation of the source IP address.¹

Facebook receives over a trillion network requests from over a billion users each day. For computational and storage reasons, we analyze different types of random 0.1% samples taken from all authenticated HTTP(S) requests to Facebook from logged-in users, as these random samples are large enough to be representative of aggregate user behavior (i.e., billions of requests from millions of distinct users). Our sampling method is deterministic over time and over network requests, selecting requests based on the hash value of a particular request attribute (e.g., user ID or IP address). As a result, our datasets include all network requests with the same randomly-selected set of attribute values over time. Specifically, we use the following network request datasets.

- **Request Random Sample:** A random sample of all network requests from Jan. 23–Apr. 19, 2020.
- **User Random Sample:** All network requests from a random sample of users, selected based on the user ID, from Jan. 23–Apr. 19, 2020.
- **IP Random Sample:** All network requests from a random sample of IP addresses, from Apr. 13–19, 2020.
- **IPv6 Prefix Random Sample:** All network requests from a random sample of IPv6 prefixes, for various IPv6 prefix lengths, from Apr. 13–19, 2020. We specifically consider IPv6 subnets of size /112, /96, /80, /76, /72, /68, /64, /60, /56, /52, /48, /44, /40, /36, and /32.

To study attacker behavior, we rely on Facebook’s measurement of abusive accounts, which are accounts created to conduct actions that violate Facebook’s policies [10], such as propagating spam or malicious content. Facebook labels such accounts through a combination of machine learning classifiers, heuristic rules, and manual investigations. Our abusive accounts dataset consists of millions of high-confidence abusive accounts labeled by Facebook as of May 3, 2020. In particular, accounts active during the time period of our study will have had at least two weeks to be caught by detection systems and labeled as abusive. We join our abusive account dataset with our network request datasets, providing the equivalent types of random samples from abusive account network requests. We note that Facebook takes actions against detected abusive accounts, affecting their behavior and our analysis. We elaborate on this impact in Section 3.3.

3.2 Ethics

While we are not directly interacting with users, our study is an empirical investigation of their IP address usage on Facebook. Users consent to Facebook’s collection and analysis of this data, per Facebook’s data policy [11]. Although IRB approval is not applicable to this research, we take care with our data collection and analysis. We only use telemetry necessary for our evaluation from random samples of logged-in Facebook users. Our study follows the basic principles of the Belmont Report [14], as it has (1) respect for persons (as logged-in users have agreed to Facebook’s terms and policies allowing for such data collection), (2) beneficence (as we limit our data collection to minimize risk and do not publish identifiable information, while gaining a broad empirical understanding

of IPv6 usage in practice), and (3) justice (insights gained from this study can help Facebook and other online services better protect users, who are the same subjects of the study).

3.3 Limitations

The data used in this study affords a large-scale evaluation of user-level IPv6 and IPv4 behavior. However, there are several limitations to our study, including:

- The data is collected from a single online platform, Facebook, and is thus representative of Facebook’s particular user and attacker population. Other platforms may observe different behavior if their users or attackers differ significantly from Facebook’s. We note that as Facebook has over 1.7 billion active users from around the world, our observations are over a broad portion of the global population, so many of our insights should generalize.
- This investigation considers a snapshot in time, and results may change in the future as IPv6 adoption increases. However, our findings can guide future directions, and many conclusions should continue to hold true.
- As an initial exploration into user-level IPv6 behavior, we consider a global perspective that analyzes observed account addresses en masse. However, we recognize that account IP behavior may differ across various ISPs and ASNs, as well as different types of networks (e.g., mobile). Our findings and recommendations may require adjustment when applied to specific networks. We leave the investigation of these facets for future work. We note however that accounting for network-specific behavior can introduce significant complexity into real-world security policies, and there is still utility in understanding aggregate behavior and identifying a single global security policy.
- Facebook’s active detection of and actions against abusive accounts influence the observed behavior of these accounts. In particular, the vast majority of observed abusive accounts are detected within a day of becoming active, indicating that our abusive account population is heavily skewed towards accounts active for short time periods (although some abusive accounts are active for longer periods). With our dataset, we are unable to determine the activity life span of individual accounts and must analyze them all together. Thus, our findings are biased towards the activity of short-lived abusive accounts. During our analysis, we will discuss how the life span skew of our abusive account population impacts our results and recommendations. We note that analyzing attacker behavior *in situ*, when defensive efforts are ongoing, is challenging. Observing the natural behavior of attackers in the absence of defenses would be insightful, but would require permitting abusive accounts to operate unchecked, an option we consider unethical and impractical. Alternatively, analyzing only accounts that maintain activity for certain lengths of time would result in a different population skew towards abusive accounts that manage to evade detection to a degree, perhaps through modifying their behavior. Ultimately, a feasible unbiased analysis would require segmenting the attacker population by their life spans and evaluating each segment individually, which we are unable to do with our dataset. We leave such investigations to future work.

¹Facebook uses IP geolocation data with high country-level precision and coverage derived from a combination of internal sources and external vendors.

- We will not be sharing the analyzed data. We recognize that this restriction does limit replication efforts, but we aim to provide generalizable insights and explain our methodology in enough detail for it to be reproduced on data from another vantage point on the internet.
- During the course of our study, the COVID-19 outbreak went from a localized event to a global pandemic. Because the pandemic impacted people’s behavior (and in particular their mobility) almost everywhere, IPv6 usage patterns may have likewise shifted. We explore these changes to the extent possible with our datasets. However, the attacker dataset and the random samples of requests based on IP addresses and prefixes were only collected later on during our study, and we cannot fully explore all dimensions of the pandemic’s impact. Thus results may differ once the pandemic’s effects subside. We note though that our investigation indicates that while there have been some understandable changes in IPv6 behavior, the changes are not drastic and our insights should hold post-pandemic.

4 DATA CHARACTERIZATION

In this section, we characterize our data demographics as well as provide an updated view of the global IPv6 landscape from Facebook’s perspective. In particular, we analyze IPv6 usage across time, ASes, and countries. Notably, we observe the impact that the global COVID-19 pandemic had on IPv6 usage, and its implications on our data analysis. We note that multiple sources [13, 19] provide similar visibility into current global IPv6 deployment levels. We consider it important to still conduct this data characterization here to provide context on our study’s datasets, which are taken from a particular vantage point at specific points in time. We relegate less relevant details to the Appendix.

4.1 IPv6 Prevalence over Time

To evaluate IPv6 prevalence over time, we consider the proportion of both users and network requests using IPv6, using the user and request random sample datasets, respectively. In Figure 1, we depict the daily proportion of Facebook users and requests that arrive from IPv6 addresses from Jan. 23–Apr. 19, 2020. On a daily basis, we observe that between 34–37% of users make some requests over IPv6, and 22–25% of requests are sent over IPv6. The IPv6 prevalence among requests is expectedly lower than among users, as we count users among the IPv6 population if they send any IPv6 requests, whereas a user’s requests are often split between IPv4 and IPv6.

We observe two significant temporal effects during our datasets’ time frame. First, we see weekend effects, where user IPv6 prevalence decreases slightly on weekends and request IPv6 prevalence increases. Second, we observe long-lasting decreases in user IPv6 prevalence and increases in request IPv6 prevalence starting in mid-March. These shifts occurred during the COVID-19 pandemic when countries around the world began locking down, and many users were quarantined at home. For example, Italy was the first European country to lock down on Mar. 9, and the first state in the US locked down on Mar. 19. Note that our datasets do not precede the first regional lockdown in China on Jan. 23, although aggregate behavior remains consistent in the early months, as Facebook has

a limited Chinese user base. We hypothesize on the causes of these temporal effects in Appendix B.

These temporal effects impact our analysis in two ways.

- (1) Due to the weekend effects, when analyzing on a day granularity, we repeat our analysis over different days. However, we generally find that the weekend effect is relatively small. Thus, we report on one day’s results and elide the analysis of additional days.
- (2) Because of the COVID-19 pandemic’s impact, when possible we perform our analysis on data both before and after the pandemic’s impact began. This comparison allows us to determine whether our observations may shift significantly once the pandemic subsides, and indicates the temporal generalizability of our results. However, we do not find significant changes in our findings. While specific statistics do differ slightly, they do not have meaningful implications. This minimal impact is not surprising in retrospect; while there are distinct changes in aggregate IPv6 behavior during the pandemic, the magnitude of the changes is ultimately limited to only a few percent. As a consequence, our statistics often differ by only a few percent as well. Thus, for clarity of exposition, we conduct the majority of our study’s analysis on data from Apr. 13–19, 2020, which is the overlapping time frame among our datasets. In Appendix A, we discuss the impact of the pandemic on individual analyses.

4.2 Prevalence by ASN

In Table 1 of Appendix A.1, we rank the top 10 ASNs by the percentage of their users using IPv6, considering ASNs with more than 1K users in the user random sample dataset taken during Apr. 13–19. We find high IPv6 deployment by these top networks, many of which provide primarily mobile network services (e.g., Reliance,




Figure 1: Daily proportion of users and requests from IPv6 addresses. The dotted vertical lines are at each Saturday. Note that the ranges of the y-axis begin at 34.5% and 22.5% for users and requests, respectively.

T-Mobile). We also note that 10.7% of the ASNs had no IPv6 usage and 28.3% had less than 10% of associated users on IPv6.

4.3 Prevalence by Country

Figure 12 in Appendix A.2 depicts a choropleth of the IPv6 user proportions across countries from Apr. 13–19, for countries with more than 1K users in the user random sample dataset. In Table 2 of Appendix A.2, we list the top countries during Jan. 23–29 and Apr. 13–19. We note that excluding Germany, the top countries remain relatively stable across the months. We discuss Germany’s changes, as well as other notable country-level changes, further in Appendix A.2.

4.4 IPv6 Client Address Patterns

Using our user random sample dataset, we explore the IPv6 client addresses themselves to understand the use of IPv4-to-IPv6 transition protocols and certain addressing schemes. We consider *6to4* [4] and *Teredo* [22], two popular IPv4-to-IPv6 transition protocols that we can identify based on client addresses. Teredo clients are assigned IPv6 addresses starting with prefix $2001::/32$, whereas 6to4 clients are in $2002::/16$. Extrapolating from our sample taken during Apr. 13–19, we observe less than 0.01% of our IPv6 user population using either of these transition protocols. We deduce that the vast majority of our users are using native IPv6, which is consistent with other data sources [13, 19]).

RFC 7707 [18] notes that some IPv6 devices embed their MAC address into the 64-bit network interface identifier (IID) portion of the IPv6 address (i.e., the final 64 bits), instead of a random value as proposed by RFC 4291 for privacy reasons [21]. The most common embedding method is to split the 48-bit MAC address in half and insert $ff:fe$ into the middle to create a 64-bit identifier. In our user random sample dataset, we observe this embedding for approximately 2.5% of IPv6 users during Apr. 13–19. Out of the subset of these users with multiple IPv6 addresses during the one-week period, 83% reuse the same IID across addresses, indicating that a static MAC is used. In the remaining cases, it is likely that the client uses MAC randomization [20], resulting in a dynamic IID. Thus, MAC-based IPv6 addresses are a non-trivial but overall small population, and most clients likely use randomized IIDs.

5 USER-CENTRIC BEHAVIOR

The effectiveness of IP-based security applications depends on the diversity and life span of benign user and attacker IP addresses, which we characterize in this section. We consider both the spatial distribution of addresses across subnets of different sizes, as well as temporal aspects about how long users use an address. We rely on the user random sample dataset for this analysis, as it provides all network requests and their source IP addresses for sampled users.

5.1 IP Addresses per User

We first consider the number of IPv4 and IPv6 addresses used by users, and how this number evolves over time.

5.1.1 RQ1 (User Behavior). Figure 2 depicts the CDFs of the number of IPv4 and IPv6 addresses observed for users over a day and over a week. (We observe similar results for other time periods.) We

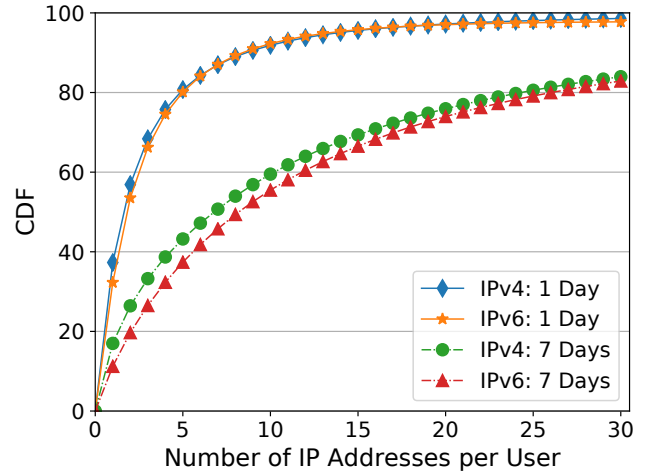


Figure 2: CDFs of the number of IPv4 and IPv6 addresses observed for users on Apr. 19 and Apr. 13–19.

see that in aggregate, users have slightly more IPv6 addresses than IPv4 addresses within the same time frame. On a single day, 32% of IPv6 users had one IPv6 address, while 20% had over 5 addresses. In comparison, 37% of IPv4 users had one IPv4 address in a day, with 19% having more than 5 addresses. Over time, the number of IP addresses per user steadily increases for both protocols, but users still have more IPv6 addresses than IPv4 addresses. During a week period, users have a median of 6 IPv4 addresses and 9 IPv6 addresses. Users likely obtain more IPv6 addresses over time as common methods for IPv6 address assignments, namely privacy-extended stateless address autoconfiguration (SLAAC) [28] and temporary-mode DHCPv6 [26], provide short-lived addresses (often with daily expirations) where new addresses have randomized IIDs. For IPv4, users may reside behind the same network address translation (NAT) and hence maintain the same public IPv4 address for longer periods. In addition, users may have multiple devices that would each be assigned distinct IPv6 addresses, but would share the same public address under IPv4.

5.1.2 RQ2 (Attacker Behavior). Figure 3 plots the CDFs of the number of IPv4 and IPv6 addresses for abusive accounts for one day. (We observe similar results for other days.) This figure shows that for both protocols, abusive accounts tend to have lower numbers of IP addresses compared to benign users (as seen in Figure 2). The majority of abusive accounts use only 1 IP address in a day for both IPv4 or IPv6. This low number is likely due to (1) quick detection by Facebook, limiting the abusive account’s activity life span (as discussed in Section 3.3), and (2) benign users organically obtaining more IP addresses as they may use their accounts on multiple devices and switch between multiple networks, even within a day. This observation may also reflect the strategies and resources employed by attackers, such as operating abusive accounts at scale using a botnet, where each abusive account is assigned to one compromised host. We elide the CDF curves for longer time periods, as our abusive account population is heavily skewed towards those active for only a day. However, we briefly note that the CDF curves are slightly lower when considering a week period, indicating that

abusive accounts that remain active over longer periods do obtain new IP addresses, rather than remain on the same ones.

We observe that abusive accounts use fewer IPv6 addresses than IPv4 addresses, which is the opposite of benign users. We hypothesize that this effect may arise because abusive accounts are sometimes forcibly cycled to new IPv4 addresses over time (even within a day) due to IPv4 address contention and NATing, whereas they may remain on the same IPv6 address (for a day) if desired. Another possible explanation is that attackers may rely on resources (e.g., compromised hosts in a botnet) with limited IPv6 capabilities.

5.1.3 RQ3 (Outlying Behavior). For both IPv4 and IPv6, we observe benign users with a large number of addresses. From the user random sample dataset for the week of Apr. 13–19, 114 users had more than 1K IPv4 addresses that week, with the largest number being 6.9K. In contrast, only 4 users had more than 1K IPv6 addresses, with a maximum of 3.5K addresses, and the IPv6 user ranked 114th had 394 addresses. We observe similar numbers for other weeks. Extrapolating from our sample, the prevalence of IPv6 outliers with over 1K addresses relative to the total IPv6 population is only 1/12 of the prevalence of IPv4 outliers. Thus, IPv6 users exhibit less extreme outlying behavior.

For attackers, we also observe fewer abusive accounts with large numbers of IPv6 addresses compared to IPv4. Considering the same week as for benign users (as some abusive accounts do remain active over this period), there are no abusive accounts in our sample with more than 1K IPv6 addresses, and only seven with more than 100 addresses (which extrapolates from our sample to several thousand abusive accounts in total across the whole platform). In contrast, 16 abusive accounts in our sample had more than 1K IPv4 addresses, with a maximum of 11K addresses. Note that the raw number of abusive accounts with many IP addresses is lower than that of benign users, likely because most abusive accounts are short-lived and have limited opportunity to obtain large numbers of addresses. We cannot conclude that abusive accounts would naturally exhibit less extreme outlying behavior than benign users if left undeterred.

A natural question is whether top users are abusive accounts missed by Facebook’s detection systems. We manually inspected the top 20 users for both IPv4 and IPv6, and did not find clear evidence that these users were abusive. We note that we are uncertain as to exactly why these users obtain so many IP addresses².

5.2 IPv6 Prefixes per User

Here we consider the diversity of user IPv6 addresses at the prefix granularity, to understand if new IPv6 addresses that users obtain primarily reside within certain prefix sizes. In essence, we are investigating the distribution of a user’s IPv6 addresses across subnets in the IP address space. Note that we do not conduct a comparison with IPv4 as IPv4 prefix behaviors are different due to different address lengths.

5.2.1 RQ1 (User Behavior). Figure 4a shows the percentage of users whose IPv6 addresses span one, two, and three prefixes during a one-week period, for different prefix sizes. We observe modal shifts at

²We do observe that the top five users reside within mobile ASNs, although the remaining top users exhibited ASN diversity.

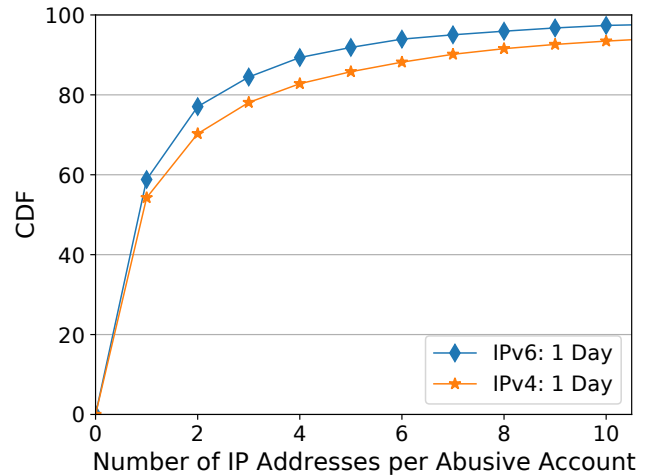


Figure 3: CDFs of the number of IPv4 and IPv6 addresses observed for abusive accounts on Apr. 19. Note that the x-axis goes up to 10 IP addresses, whereas in Figure 2 it goes up to 30.

/64 and prefixes shorter than /48, indicating significant aggregations of user IPv6 addresses within prefixes of these lengths.

For prefixes longer than /64, the distribution of the number of prefixes per user is similar to the distribution of IP addresses per user (shown in the graph as a /128 prefix). This property indicates that for many users, their IP addresses tend not to share prefixes longer than a /64. Thus when many users are assigned a new IPv6 address, it is usually a different /68 within the same /64. This address assignment pattern matches the behavior of using 64-bit address IIDs (following RFC recommendations [3, 21]) with randomization (such as provided by privacy-extended SLAAC [28]). Note that this pattern is in contrast with common server address assignments, where low-order bytes of the address are varied, resulting in multiple servers sharing the same long prefix (e.g., a /112) [27].

We also observe additional aggregations at prefixes shorter than a /48. RFC 4291 [21] specifies the format of global unicast IPv6 addresses to contain a global routing prefix, followed by a subnet ID then the network interface ID (IID). The IID is typically 64 bits [3, 21] (also discussed in Section 4.4). These aggregating prefixes are likely the global routing prefixes for the networks that users are on. Thus, there are many users with addresses spanning multiple /64 subnets, rather than being constrained within a single /64 subnet, but their addresses aggregate within the network’s routing prefix.

5.2.2 RQ2 (Attacker Behavior). For abusive accounts in Figure 4b, we still observe a modal shift at /64 and for prefixes shorter than a /48, similar to benign users. These shifts indicate that the abusive accounts with multiple IP addresses tend to still reside within the same /64 subnet, and those with more addresses across multiple /64 subnets further aggregate within their network’s global routing prefix. Note that as our abusive account population is skewed towards short-lived accounts (with few IP addresses to begin with), we cannot directly compare the prefix diversity of abusive accounts and users over a week period.

5.2.3 RQ3 (Outlying Behavior). The addresses of users with a large number of IPv6 addresses tend to be spread out across different

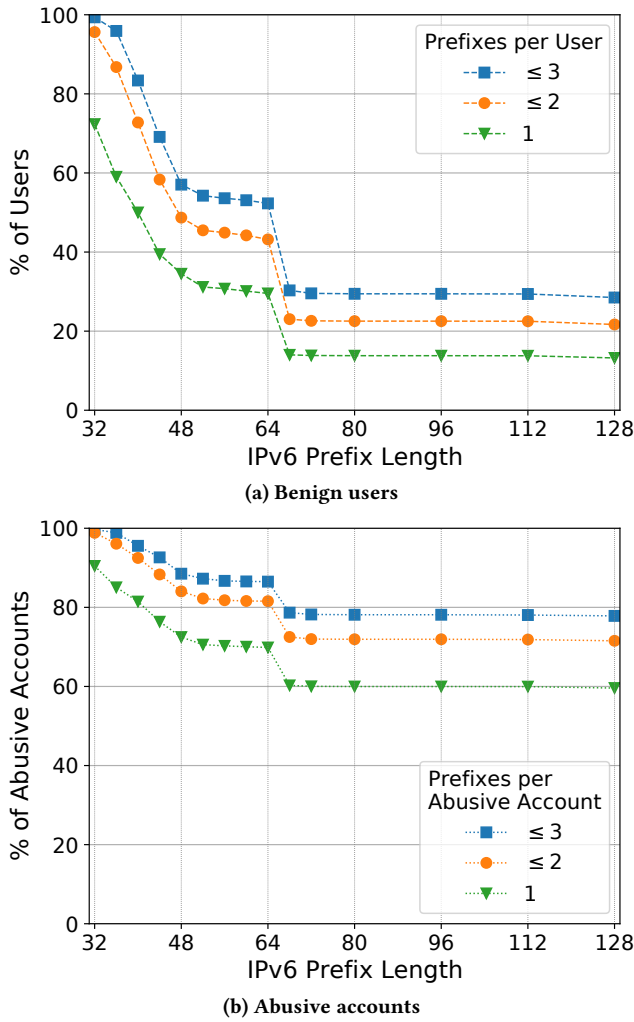


Figure 4: The percentage of benign users and abusive accounts whose IPv6 addresses span 1, 2, and 3 prefixes during a week period (Apr. 13–19), for varying prefix lengths. Both populations exhibit aggregation of addresses within subnets of similar prefix sizes (e.g., /64). Note that the population proportions for abusive accounts and benign users *should not* be compared here, as our abusive account population exhibits skew in its activity life span.

/64 subnets. From Section 5.1.3, the 4 users in the user sample dataset with more than 1K IPv6 addresses also had more than 1K /64 prefixes. Overall there are 58 users in our user sample with more than 500 IPv6 addresses and 39 users with more than 500 /64 prefixes. However, we begin observing aggregations at the /44 level, where only 12 users had more than 100 /44 prefixes, with a maximum of 143 prefixes. Again, this aggregation is likely within the global routing prefix of a single network. Thus, users with high numbers of IP addresses also exhibit high subnet diversity within a network. Attackers exhibit the same pattern, although with less prefix diversity, as there were fewer extreme outliers in terms of IPv6 addresses per abusive account (from Section 5.1.3).

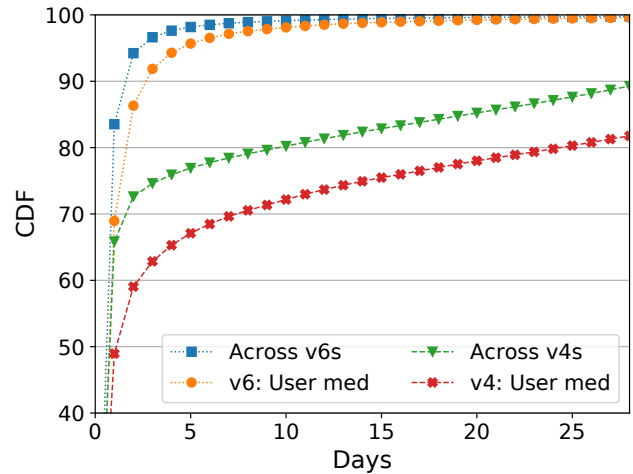


Figure 5: CDFs of the number of days since we first observed an IP address for a user (i.e., the address’s life span for the user), for both IPv4 and IPv6 addresses observed on Apr. 19. We plot the life span distributions across all (user, address) pairs as well as the median per user. (Note that the y-axis begins at 40%.)

5.3 IP Life Spans for Users

Finally, we consider the longevity of user IP addresses.

5.3.1 RQ1 (User Behavior). For each (user, IP address) pair that we observed on Apr. 19, Figure 5 depicts the CDFs of the number of days since we first observed that pair, which represents an IP address’s life span for a user. We plot the life span distributions across all pairs, as well as taking the median per user (aggregating across the addresses observed for each user). We see that IPv6 addresses are far shorter-lived for users than IPv4 addresses, likely due to frequent expirations of IPv6 address assignments [26, 28]. While 66% of (user, IPv4 address) pairs had not been seen in any of the 27 days prior to Apr. 19, 84% of (user, IPv6 address) pairs had this property. Only 1.2% of IPv6 pairs had been first observed more than a week prior to Apr. 19, compared with 22% of IPv4 pairs. When grouping the addresses per user, we observe that the CDF of the median IP address life span per user is lower than across all pairs, for both IPv4 and IPv6. This property indicates that while most users’ IP addresses (particularly for IPv6) are short-lived, users do maintain activity on some IP addresses for longer periods.

In Section 5.2 we observed that user IPv6 addresses are often aggregated within /64 prefixes and prefixes shorter than a /48. If so, one might expect that while a user’s IPv6 addresses are short-lived, they may remain active within IPv6 prefixes for longer. We can explore this hypothesis using Figure 6a, which plots the percentage of (user, IP prefix) pairs that were first observed within the past 1, 2, and 3 days (i.e., the prefix’s life span for the user), across all pairs for both IPv6 and IPv4. As expected for IPv6, we observe modal shifts around the /64 and /48 levels, where the percentage of short-lived (user, prefix) pairs drops drastically. These decreases indicate that users are longer-lived within /64 prefixes, and even longer-lived within prefixes shorter than a /48 (likely the global routing prefixes for networks). We note that the life span distribution for IPv4 addresses is most similar to that of the IPv6 /64 prefix for users.

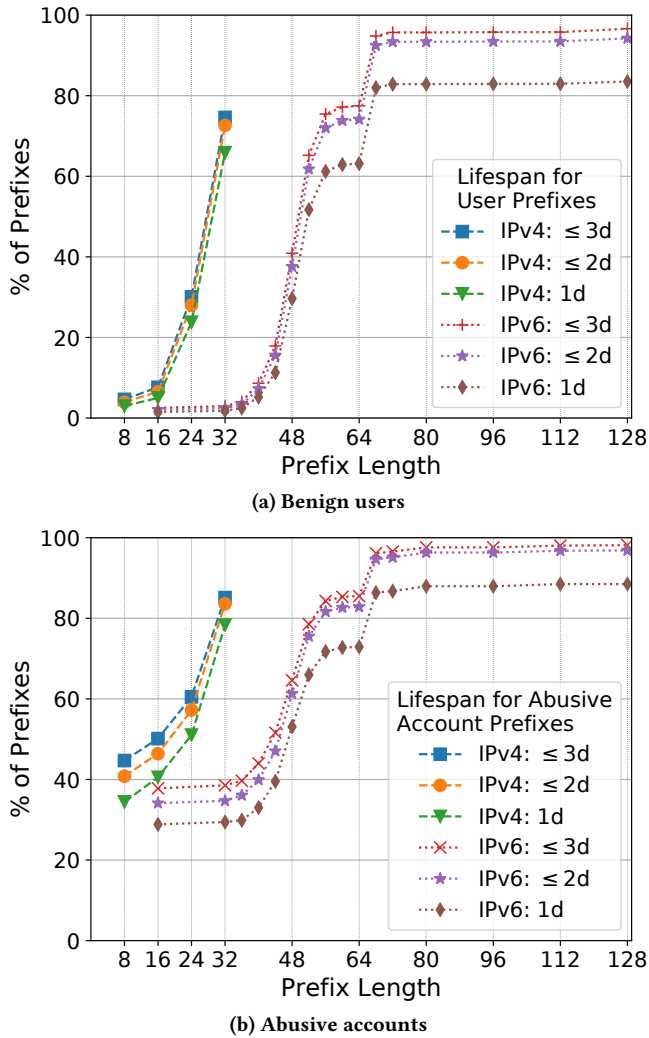


Figure 6: For various-length IP prefixes corresponding to both benign user and abusive account IP addresses, we depict the percentage of IP prefixes whose user life spans (i.e., days since a (user, prefix) pair was first observed) are 1, 2, and 3 days, for both IPv4 and IPv6. As before, benign users and abusive accounts *should not* be directly compared due to the prevalence of short-lived abusive accounts.

Prior work [29, 37] studied the dynamics of IPv4 addresses, observing that a significant portion of addresses change within 1–3 days, often across /24 prefix boundaries. Our IPv4 results concur with these previous observations.

5.3.2 RQ2 (Attacker Behavior). In Figure 6b, we plot the percentage of (abusive account, IP prefix) pairs where the pair was first observed within the last 1, 2, or 3 days, across all pairs for both IPv6 and IPv4. We observe a similar pattern as with benign users, where abusive accounts (that do remain active for longer periods) reside longer in /64 prefixes and prefixes shorter than a /48, compared to on IPv6 addresses. As before, we cannot directly compare benign users and abusive accounts due to the life span skew of our abusive account population.

5.3.3 RQ3 (Outlying Behavior). Benign users do use IP addresses for long periods of time (exceeding 28 days), but these outlier address life spans are more prevalent for IPv4 than for IPv6. Only 0.23% of (user, IPv6 address) pairs exceed 28 days, compared with 10.7% of (user, IPv4 address) pairs. We were unable to determine exactly why users reside on these addresses for long periods of time. Users may be on devices with statically allocated addresses for either protocol, or for IPv4, users may remain behind the same NAT with the same public-facing IP address. We note that these addresses are distributed broadly among many ASNs and countries. We observe few outliers for attackers, although this observation is rather an artifact of Facebook’s defenses.

6 IP-CENTRIC BEHAVIOR

In this section, we take an IP-centric view of internet traffic, evaluating the user populations per IP address and prefix, and comparing IPv6 with IPv4. One could also consider how long users reside on an IP address and prefix; this analysis is actually already provided by Section 5.3, where the user-centric analysis of how long an address or prefix is used by a user is equivalent. This IP-centric perspective is important as defense mechanisms such as IP blocklisting or rate limiting enforce at the IP level, impacting all users on an IP address or prefix. Thus to assess the potential impact of such actions, it is important that we develop an understanding of how users populate IP addresses.

Recall that our abusive account population is skewed towards short-lived accounts, due to Facebook’s detection technology and consequent actions. Attackers likely modify their behavior in response to Facebook’s actions, such as creating more abusive accounts to continue their attacks. The abusive account populations identified per IP address and prefix over time reflect this reality.

6.1 Users per IP Address

We first investigate the number of users per IP address for both IPv4 and IPv6. To conduct this analysis, we use the IP random sample dataset, which contains all network requests (from all users) for a random sample of IP addresses.

6.1.1 RQ1 (User Behavior). Figure 7 depicts the CDFs of the number of users per IPv4 and IPv6 address, both for one day and one week. We observe that IPv4 addresses exhibit significantly more users per address than IPv6, which is expected as IPv4 addresses are often NAT’ed to deal with IPv4 address exhaustion. Only a third of IPv4 addresses had a single user on it during a day, compared with 95% of IPv6 addresses. In fact, over 99% of IPv6 addresses had only one or two users, and in many cases the second user may simply be sharing the same device with the first user.

Taking a longer view, the number of users per address increases considerably over time for IPv4, but not for IPv6. For example, the percentage of single-user IPv4 addresses decreases to 23% after a week, while for IPv6, it decreases by only 0.1%.

6.1.2 RQ2 (Attacker Behavior). For IP addresses with abusive accounts, we inspect the number of abusive accounts per address, as well as the number of benign users per address. Figure 8 depicts the distribution of these per-IP populations for IP addresses with




Figure 7: CDFs of the number of users observed per IP address over a day (Apr. 13) and a week (Apr. 13–19), for both IPv4 and IPv6.

at least one abusive account observed on it, for both a day and a week period.

By comparing with Figure 7, we observe that overall there are fewer abusive accounts per IPv4 address than benign users. Whereas a third of IPv4 addresses only had a single user on it in a day, 73% of IPv4 addresses with an abusive account had only one. For IPv6, benign user and abusive account behaviors were similar: out of all IPv6 addresses with at least one abusive account, about 95% of addresses had a single abusive account. Thus, attackers tend not to use a large number of abusive accounts on a single IP address. This pattern does not shift significantly over longer time windows (a day versus a week), suggesting that attackers typically are not reusing IP addresses when creating new abusive accounts to replace detected ones.

We do observe a notable difference between IPv4 and IPv6 in the number of benign users on IP addresses with abusive accounts. For IPv6, 63% of addresses only had abusive accounts and no benign users in a day, and only 12% had more than one benign user. In comparison, only 3.4% of IPv4 addresses had only abusive accounts in a day, and 72.9% of addresses with abusive accounts had more than 10 benign users. Thus, abusive accounts tend to be more isolated on IPv6 addresses, but share IPv4 addresses common to many users. This difference has significant implications for the collateral damage that may be imparted when taking IP-level actions for IPv4 versus IPv6 addresses.

6.1.3 RQ3 (Outlying Behavior). For both IPv4 and IPv6, we observe addresses with large user populations. In the IP random sample dataset, 14 IPv4 addresses had over 200K users in a week, including an address with 830K users and another with 750K users. For IPv6, we observe less extreme behavior, with at most 71K users for an IPv6 address in a week, and the next highest number being 30K users. Beyond these two, only 10 IPv6 addresses in our sample had over 1K users, and only 22 had over 150 users. Meanwhile, there were 4402 IPv4 addresses with over 1K users. Extrapolating from our IP random sample, the proportion of IPv6 addresses with more

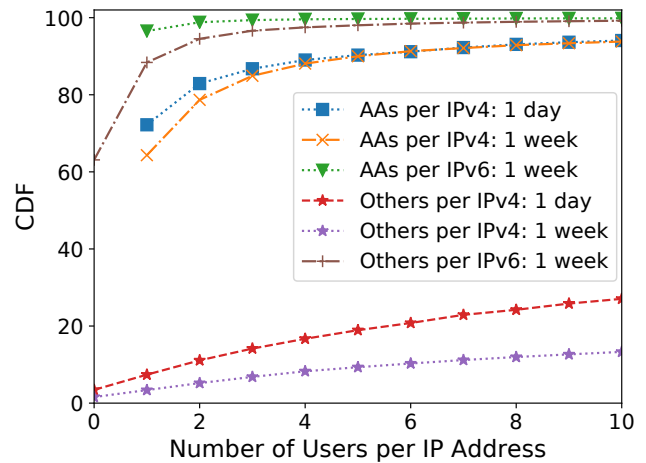


Figure 8: CDFs of the number of abusive accounts (AA) observed per IP address over a day (Apr. 13) and a week (Apr. 13–19), for both IPv4 and IPv6 addresses with at least one abusive account observed. We also depict the CDFs of the number of benign users observed for these same IP addresses during the same time period. We elide showing the single-day IPv6 curves as they resemble those for a week.

than 1K users is more than four orders of magnitude (12,720 times) smaller than the proportion for IPv4.

Our IP random sample dataset provides accurate user counts per address, but may have missed addresses with more extreme behaviors. However, we can further investigate the population of addresses with many users by using our user random sample dataset. As this dataset provides all network requests (and their IP source addresses) for a random 0.1% sample of users, addresses that have more than 10 users in the user sample have in expectation more than 10K users in the full dataset. For the same week (Apr. 13–19), we observe 7887 IPv6 addresses with more than 10 sampled users, with at most 89 sampled users (corresponding to an expected 89K users). In comparison, we observe 659K IPv4 addresses with over 10 sampled users, with a maximum of 3.9K users (indicating millions of users on that IPv4 address).

Our IP and user random sample datasets paint a consistent picture. Hundreds of thousands of IPv4 addresses have massive numbers of users (up to millions of users per address). In contrast, several thousand IPv6 addresses have more than 10K users, with none exceeding 100K users.

We also consider the ASNs that assign these heavily populated IP addresses. Using the user random sample dataset, we observe that for IPv6, 96% of IPv6 addresses with more than 10K users per address (in expectation) belong to ASN 20057 (AT&T Mobility). Diving deeper, we observe that these heavily populated ASN 20057 IPv6 addresses have a distinct address structure³ compared with more sparsely populated IPv6 addresses from that ASN, making creating signatures for heavily populated IP addresses feasible. Beyond this ASN, ASNs 13335 (Cloudflare), 16276 (OVH) and 14061 (Digital Ocean) account for another 2.8% of IPv6 addresses with

³The IPv6 addresses of ASN 20057 use IID randomization except for the detected heavily populated IPv6 addresses, where the IID bits are all zeros except for the least significant 16 bits.

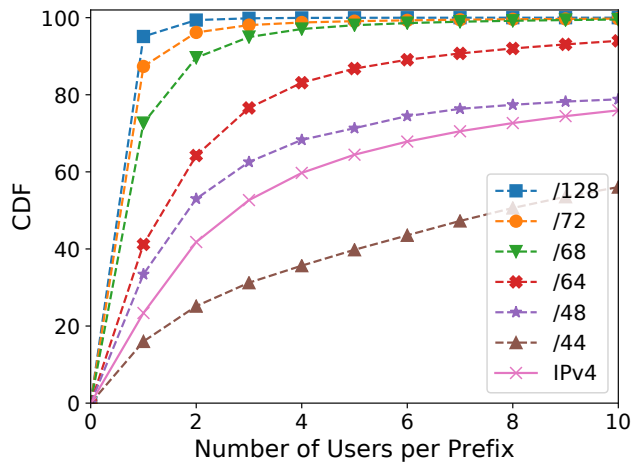


Figure 9: CDFs of the number of users observed within IPv6 prefixes of varying sizes in a week (Apr. 13–19).

more than 10K users, and a short tail of 13 other ASNs account for the remaining 1%. For IPv4, there is a much larger set of 1568 ASNs with heavily populated IP addresses. ASN 23693 (Telkom Indonesia) accounts for the largest proportion at 22%, with the next three ASNs, 24203 (Axiata), 4761 (Indosat), and 38266 (Vodafone India), accounting for an additional 10%. Thus, for IPv4, heavily populated IP addresses are much more prevalent in sheer number as well as across networks.

For attackers, we observe outlying behavior primarily for IPv4. In our IP random sample, 16 IPv4 addresses had more than 1K abusive accounts, with a maximum of 10.8K abusive accounts. For IPv6, no sampled IP address had over 1K abusive accounts and only 7 had more than 100 abusive accounts. Extrapolating to the total population implies that at least 10K IPv4 addresses had more than 1K abusive accounts, and the number of IPv6 addresses with more than 1K abusive accounts is in the hundreds at most. We note that for both IPv4 and IPv6, the heavily populated IP addresses had abusive accounts on them, and addresses with many abusive accounts also had many benign users. Thus, collateral damage may be high if acting on these IP addresses.

6.2 Users per IPv6 Prefix

We now consider the number of users per IPv6 prefix, for varying prefix sizes. We do not compare with IPv4 prefixes, which differ due to different address lengths.

6.2.1 RQ1 (User Behavior). Figure 9 depicts the CDFs of the number of users per prefix in a week, for varying prefix sizes. We observe that for /72 prefixes and longer (elided from the figure), the distributions are similar to that of full IPv6 addresses (represented as the /128 prefix). For example, 87% of /72 prefixes have a single user, compared with 95% of IPv6 addresses. However, the distributions shift significantly when considering shorter prefixes, indicating that there is an aggregation of users per prefix, starting at the /68 level. The largest shift is with /64 prefixes, where only 41% of /64 subnets contain a single user, as compared with 73% of /68 subnets. This result indicates that groups of users (e.g., home or enterprise networks) are often allocated within /64 subnets. Another sizable


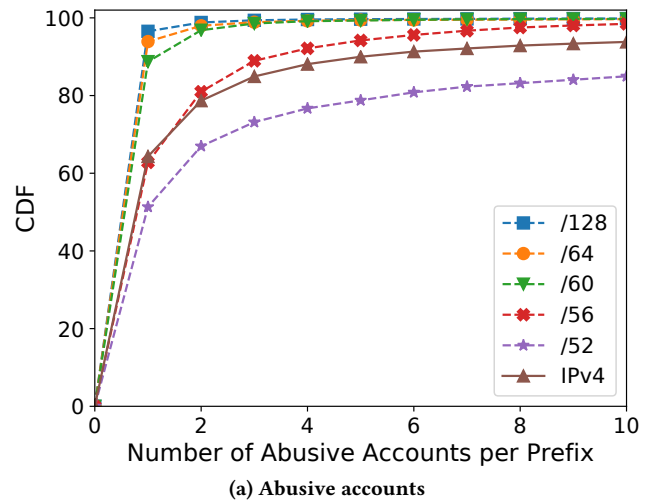


Figure 10: For IPv6 prefixes of varying sizes containing abusive accounts, we depict the CDFs of the number of abusive accounts and benign users observed within each prefix in a week (Apr. 13–19).

shift occurs between the /48 and /44 prefixes, where the percentage of single-user prefixes decreases from 34% to 16%. This shift likely represents the aggregation of users within a network’s global routing prefix. We note that in this dimension, IPv4 addresses appear most similar to IPv6 /48 prefixes.

Recall our observation from Section 5.2 that IP addresses for the same user aggregate within /64 or larger subnets. Here, we similarly observe significant aggregation of users within /64 subnets, but we begin seeing user aggregation within the same /68 prefixes. We hypothesize that this /68-level aggregation arises largely due to randomness, where address IID randomization [26, 28] by multiple users’ clients within the same /64 subnet results in some users within the same /68 subnet by chance⁴.

⁴A user with three IPv6 addresses assigned using IID randomization has about a 20% chance of being in the same /68 subnet as another user’s IP address within the same /64 subnet. This probability increases with the number of IP addresses obtained by either user, which naturally increases over time (Section 5.1).

6.2.2 RQ2 (Attacker Behavior). Figure 10a depicts the number of abusive accounts per prefix in a week, for varying prefix lengths. Unlike the trend for benign users, there tends to be little aggregation of abusive accounts even at /64 prefixes. Only at the /56 level do we see a significant increase in the number of abusive accounts per prefix. This result suggests that abusive accounts tend to be distributed more broadly across prefixes than users, spanning many /64 subnets rather than aggregating within a single one.

From Figure 10b, which plots the CDFs of the number of benign users in prefixes with abusive accounts, we observe that prefixes with abusive accounts tend to also be heavily occupied by benign users. Comparing with Figure 9, we can see that prefixes with abusive accounts have more users than the overall distribution among prefixes. For example, while 41% of all /64 prefixes contain a single user, only 19% of /64 prefixes with abusive accounts contain zero or one benign users. This pattern is similar to the behavior of abusive accounts on IPv4 addresses studied in Section 6.1, where we found that abusive accounts tend to use IPv4 addresses heavily populated by benign users. Overall, /56 subnets with abusive accounts appear most similar to IPv4 addresses in their abusive account vs. benign user population distributions.

6.2.3 RQ3 (Outlying Behavior). Outlying behavior for IPv6 prefixes is similar to that of full IPv6 addresses. As there are single addresses with more than 10K users, there are also prefixes with a large number of users. As in our analysis of users per IP address, we can use the user random sample dataset to determine the expected number of users per prefix, across all prefixes. Interestingly, while we observe that users typically aggregate at the /64 level, we observe heavily populated IPv6 addresses aggregating even at the /112 level. While we observed at most 89K users per IPv6 address, we observe a /112 prefix with 2.3M users. There are 39 /112 prefixes with more than 1M users. At the /64 level, we still only observe 39 /64 subnets with more than 1M users, indicating these /112 subnets dominate the outlying users per prefix behavior. We find that all of these /112 subnets belong to ASN 20057 (AT&T Mobility), which also accounted for 96% of IPv6 addresses with large user populations (from Section 6.1).

In total, we observe 393 /64 subnets exceeding 10K users (in expectation), distributed across 27 ASNs. ASN 9009 (M247) accounted for the largest percentage of these heavily populated prefixes at 21%, with the top 4 accounting for 61% of these prefixes. Recall from Section 6.1 that 17 ASNs contained 7.9K IPv6 addresses with more than 10K users. Thus, the majority of heavily populated IPv6 addresses aggregate within a small set of /64 subnets, and heavily populated IPv6 prefixes are limited in quantity and AS diversity. We note though that these prefixes do often contain abusive accounts, and thus actioning on them could have high collateral impact.

Similar to IPv6 addresses, we observed few IPv6 prefix outliers in terms of abusive accounts, with only two /64 subnets and 13 /48 prefixes in our IPv6 prefix random sample dataset containing over 100 abusive accounts. Again, these prefixes with many abusive accounts also had many benign users.

7 IPV6 IN SECURITY SYSTEMS

In Sections 5 and 6, we characterized user and attacker behavior across IPv4 and IPv6, providing answers to our first three research

questions. In this section, we draw on our insights to tackle our final research question (RQ4) on how security applications should employ IPv6 addresses, particularly focusing on the use of IPv6 addresses in blocklisting, rate limiting, threat exchanges, and machine learning models.

Our analysis and discussion in this section, as in the rest of this paper, consider a purely global perspective. Conclusions may need to be adjusted if applied to specific ISPs or ASNs, particularly regarding the impact of operating at different prefix granularities. Per-ISP/ASN security policies may be more specific and accurate, and we leave investigating these differences for future work. We note though that accounting for network-specific behavior can introduce significant complexity into real-world defenses, and there is still utility in understanding aggregate behavior and identifying a single security policy applicable to all traffic.

7.1 Effectiveness of IPv6-Based Actions

We start by exploring the effectiveness of IPv6-based security actions, given the spatial and temporal characteristics of benign user and attacker IPv6 behavior. We simulate the following scenario: we count the proportion of abusive accounts per IP prefix on day n , and consider what would happen on day $n + 1$ if we actioned on all prefixes with a proportion over some threshold t . (For now, we consider all actioning generically; in Section 7.2 we explore implications for different types of actions.) On one extreme, a threshold of 0% indicates that any prefix with a single abusive account will be actioned on, which could potentially inflict significant collateral damage on benign users of the prefix. On the other extreme, a 100% threshold indicates actioning only if we expect no collateral damage, but this restraint may cause us to miss actioning on many abusive accounts on mixed-population IP prefixes.

Note that, as with our analysis in Section 6, the abusive account populations identified per prefix reflect attacker behavior in the presence of Facebook detection and actions. We consider it valuable to analyze abusive accounts under existing defensive efforts, as it provides insights on how to further combat attackers. We leave the exploration of unchecked abusive accounts for future consideration.

Figure 11 shows the Receiver Operating Characteristic (ROC) curves for this scenario for IPv6 prefixes of different sizes. We observe that when actioning on IPv6 addresses (represented as a /128 prefix), the true positive rate (TPR) is at most 14.3%. Abusive accounts frequently appear on new IP addresses, consistent with the short-lived nature of abusive accounts and our observation that abusive accounts tend not to cluster together on the same IP address (from Section 6.1.2). At a 0% threshold for taking action, the false-positive rate (FPR) is 0.9%, representing many millions of accounts on a service as large as Facebook. There are a relatively small number of abusive accounts residing on the heavily populated IPv6 addresses discussed in Section 6.1.3, resulting in high collateral impact. As a result, raising the threshold to 10% significantly decreases the FPR to 0.01%, while the TPR only drops to 13.0%, as one avoids taking action on the heavily populated IPv6 addresses. At a 100% threshold, we observe a TPR of 7.6% with a FPR of 0.0009%. Here, the FPR does not reach zero as some benign users may appear on day $n + 1$ on the IPv6 addresses that were purely malicious on day n .

In Section 5.2, we observed that many abusive account IPv6 addresses tended to aggregate within /64 subnets. Actioning at the /64 granularity should then result in higher TPRs, as some abusive accounts will appear the next day on new IPv6 addresses, but within the same /64 prefix as the previous day. However, the FPR may increase if benign users frequently occupy the same prefix as abusive accounts. At a 0% actioning threshold, we observe a TPR of 21.2% with a FPR of 0.9%, which is better than the 0% threshold for actioning on full IPv6 addresses. This improvement occurs because the FPR is dominated in both the /64 and /128 cases by the collateral impact of actioning on heavily populated IP addresses, but using /64 subnets captures more next-day abusive accounts. At a 10% threshold, the TPR of actioning on /64 subnets remains at 19.8%, while the FPR is 0.2%. However, at the 100% threshold, we observe a TPR of only 4.2% with a FPR of 0.01%. This result is strictly worse than with the full IPv6 address, as the /64 subnets that are purely abusive accounts one day often have benign users the next day, and abusive accounts in mixed-population /64 subnets are missed. Thus, we find that the optimal granularity for IPv6 actioning could be either the full address or the /64 prefix, depending on one’s false positive tolerance.

For IPv4 addresses, a 0% threshold results in a 65.8% TPR but a 27.1% FPR. Thus, IPv4 addresses with abusive accounts on one day tend to have abusive accounts the next day, perhaps due to challenges in obtaining IPv4 address diversity. However, with many benign users behind the same IPv4 addresses (as observed in Section 6.1.2), particularly due to NATing, the collateral damage is massive for actioning on IPv4 addresses. At a 10% threshold, the TPR is 9.7% with a 0.16% FPR. Overall, for FPR values below 1%, IPv4’s ROC curve is consistently below those of IPv6 addresses and /64 prefixes, indicating that at low FPR levels, IPv6 actions can be more effective than IPv4 ones. We note that the ROC curves for IPv4 addresses and /56 IPv6 prefixes are similar, which is consistent with our finding in Section 6.2 that IPv4 addresses and /56 IPv6 prefixes behave similarly when containing abusive accounts.

7.2 Implications for Defense Mechanisms

We now combine the above analysis of false positive/false negative tradeoffs with our previous insights into benign and malicious IPv6 behavior to analyze the impact of IPv6 actioning using different security mechanisms.

Blocklisting. IP blocklisting is widely used in practice (e.g., Spamhaus [34], Spamcop [33]). For IPv4 addresses, blocklisting can have significant collateral damage, as many users reside on the same IPv4 address due to NATing. Exacerbating the risk, we observed that attackers tended to operate abusive accounts on IPv4 addresses with even more users than the typical address (Section 6.1). However, IPv6 differs in that there are few users per address, and we do not observe a similar trend where IPv6 addresses with abusive accounts have more benign users than typical (likely due to the dearth of such addresses). Additionally, we observed that IP addresses with high numbers of users are limited in quantity and diversity in IPv6, primarily residing within a small set of ASNs (Section 6.1). Compared to heavily populated IPv4 addresses, these heavily populated IPv6 addresses can be feasibly predicted to avoid blocklisting and to handle through other means. Thus, IPv6 address




Figure 11: The Receiver Operating Characteristic (ROC) curves for one day when actioning on IPv6 subnets with abusive accounts the day before. We explore different prefix sizes (eliding the curves for less optimal ones) and various thresholds for actioning based on the abusive account proportion per prefix.

blocklisting can potentially be more aggressive without suffering the same collateral damage risks of IPv4.

IPv6 does introduce a challenge to blocklisting: addresses are much shorter-lived than with IPv4 for both benign users and abusive accounts (Section 5.3), and thus IPv6 blocklisting is likely most effective when deployed short term. We witnessed the impact of this temporal facet in Section 7.1, where we saw that actioning on IPv6 addresses on one day results in limited (but non-trivial) recall the next day. On the other hand, as benign and malicious IP addresses exhibited some spatial locality for subnets of size /64 or larger (Section 5.2), blocklisting on IPv6 prefixes may be more robust to temporal dynamics, although at the risk of more false positives. We determined that for low false positive rates that may be tolerated in practical situations (e.g., 0.1%), acting on /64 prefixes results in higher recall than acting on the full IPv6 address. Overall, the number of users per /64 prefix remains lower than for IPv4 addresses, in aggregate as well as when restricting to those with abusive accounts, and heavily populated /64 subnets are infrequent and more readily predicted. Generally, across various operating points, IPv6 provides better blocklisting tradeoffs than IPv4.

We also note that we found in Sections 6.2 and 7.1 that IPv4 addresses with abusive accounts are most characteristically similar to /56 subnets with abusive accounts, so existing IPv4 blocklisting policies may be applicable to /56 prefixes.

Rate Limiting. IP-based rate limiting is another popular security mechanism that shares many of the same characteristics and considerations as blocklisting, as discussed above. However, as IP-based rate limiting is naturally defined over short time windows and agnostic to the specific address, it should be less affected by temporal dynamics (such as IPv6’s short-lived addresses) than blocklisting is. While rate limiting is often employed at the user level, IP-based rate limiting has value as a safeguard against undetected attackers,

attackers with multiple abusive accounts per address, and attacks that don't require accounts (e.g., public data scraping).

For IPv4, rate limiting thresholds can be difficult to determine, as the number of true benign users per address can vary widely. This uncertainty often necessitates liberal thresholds permitting a large quantify of activity. However, for IPv6, the number of benign users per address is significantly more limited. We observed in Section 6.1 that less than 0.2% of IPv6 addresses have more than three users in a day, compared to 29.3% of IPv4 addresses. Even if applying rate limiting at IPv6 subnet granularities, we observed fewer users per /64 IPv6 prefix compared to IPv4 addresses (from Section 6.2). Thus, IPv6 rate limiting thresholds can be set more tightly, even for logged-out requests, by assuming a small number of legitimate users per IPv6 address or prefix. Additionally, as discussed with respect to blocklisting, IPv6 outliers in the number of users per address or prefix are fewer and more easily predicted for different treatment than IPv4 outliers. We found in Section 6.2 that IPv6 /48 subnets appeared most similar to IPv4 addresses in their distribution of user population sizes, so existing rate limiting logic applied to IPv4 addresses could potentially be translated to IPv6 /48 prefixes.

Threat Exchanges. Threat exchanges and threat intelligence aggregators are platforms through which online services can share information about malicious actors, in hopes that the collective intelligence will enable better security responses. In this work, we consider one class of attackers, those that operate abusive accounts, on a particular online platform. We do not consider the extent to which these attackers also target other platforms, although prior work [24, 35] has investigated these considerations. Our analysis in Section 7.1 indicates that taking action against IPv6 addresses and prefixes with abusive accounts on one day has a noticeable but limited impact on the next day. This finding suggests that the value of intelligence on suspicious IPv6 addresses degrades quickly.

Machine Learning Models. Machine learning models can use IP behavior features to classify users and IP addresses, and prior work has used IPv4 features [2, 38]. As our study found various differences in how users use IPv4 and IPv6 addresses, as well as the population of users that appear on IP addresses, models may perform better if treating the two protocols distinctly. In particular, these models could be designed to be robust against temporal dynamics, such as the shorter life span of IPv6 addresses for users (Section 5.3). The low number of users per IPv6 address may provide cleaner signals for use cases such as IP address reputation models. However, combining users within IPv6 prefixes, such as the /64 subnets that users often aggregate within (Sections 5.2 and 6.2), may provide more data for decision making, such as when clustering user groups that behave similarly, whereas IPv6 addresses are unlikely to have sizable user groups compared to prefixes. In Sections 6.2 and 7.1, we observed that attackers used abusive accounts similarly on IPv4 addresses as /56 subnets, so attack detection models may be able to employ a single set of IP features if extracted at those two granularities.

8 CONCLUSION

In this study, we analyzed IPv6 behavior at the user level, comparing with IPv4 behavior. From one of the largest online platform's

vantage point, we investigated the addresses obtained by both benign users and attackers, as well as the user populations across addresses and prefixes. We found that unlike with IPv4, IPv6 addresses are sparsely populated by users, even if attackers operate abusive accounts on them. However, they are short-lived for users, although users stay within larger (e.g., /64) subnets for longer. Our findings provide insights on deploying security mechanisms for IPv6 users, potentially performing more effectively than on IPv4.

Future work can expand upon our initial exploration of user-level IPv6 behavior. One direction is in characterizing IPv6 behavior across different network types, such as mobile, residential, and enterprise networks, as well as across different ISPs and ASNs. Mobile networks are particularly interesting given their high IPv6 deployment rates and distinct host types (e.g., smartphones). A related direction is to investigate the causes of dynamic IPv6 behavior, similar to the exploration of IPv4 dynamic address reasons by Padmanabhan et al. [29]. With such an understanding, security mechanisms could be more effectively deployed by factoring in a user's specific network and type. Further research can also evaluate IPv4-based machine learning models on IPv6, and experiment with IPv6-aware models. For example, existing malicious account classifiers [2, 38] use IPv4 features but may perform differently for IPv6 users. Finally, our study only considered one attacker type. Additional investigation into the IPv6 behavior of other attacker classes, such as those hijacking accounts and launching attacks without accounts (e.g., public data scraping), would be fruitful.

REFERENCES

- [1] Robert Beverly, Ramakrishnan Durairajan, David Plonka, and Justin P. Rohrer. 2018. In the IP of the Beholder: Strategies for Active IPv6 Topology Discovery. In *ACM Internet Measurement Conference (IMC)*.
- [2] Qiang Cao, Xiaowei Yang, Jieqi Yu, and Christopher Palow. 2014. Uncovering Large Groups of Active Malicious Accounts in Online Social Networks. In *ACM Conference on Computer and Communications Security (CCS)*.
- [3] Brian Carpenter and Sheng Jiang. 2014. RFC 7136: Significance of IPv6 Interface Identifiers. IETF Request for Comments.
- [4] Brian Carpenter and Keith Moore. 2001. RFC 3056: Connection of IPv6 Domains via IPv4 Clouds. IETF Request for Comments.
- [5] Catalin Cimpanu. 2019. Belarus becomes first country to make IPv6 mandatory for ISPs. <https://www.zdnet.com/article/belarus-becomes-first-country-to-make-ipv6-mandatory-for-isps/>
- [6] Lorenzo Colitti, Steinar H. Gunderson, Erik Kline, and Tiziana Refice. 2010. Evaluating IPv6 adoption in the Internet. In *International Conference on Passive and Active Network Measurement (PAM)*.
- [7] Jakub Czyz, Kyle Lady, Sam Miller, Michael Bailey, Michael Kallitsis, and Manish Karir. 2013. Understanding IPv6 Internet Background Radiation. In *ACM Internet Measurement Conference (IMC)*.
- [8] Stephen E. Deering and Robert M. Hinden. 1998. RFC 2460: Internet Protocol, Version 6 (IPv6) Specification. IETF Request for Comments.
- [9] Amogh Dhamdhere, Matthew Luckie, Bradley Huffaker, KC Claffy, Ahmed Elmokashfi, and Emile Aben. 2012. Measuring the Deployment of IPv6: Topology, Routing and Performance. In *ACM Internet Measurement Conference (IMC)*.
- [10] Facebook. 2020. Community Standards. <https://www.facebook.com/communitystandards/>
- [11] Facebook. 2020. Data Policy. <https://www.facebook.com/about/privacy/>
- [12] Facebook. 2020. Facebook Q2 2020 Earnings. <https://investor.fb.com/investor-events/event-details/2020/Facebook-Q2-2020-Earnings-/default.aspx>
- [13] Facebook. 2020. IPv6. <https://www.facebook.com/ipv6/>
- [14] National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. 1978. The Belmont report: Ethical principles and guidelines for the protection of human subjects of research. <https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/>
- [15] Pawel Foremski, David Plonka, and Arthur Berger. 2016. Entropy/IP: Uncovering Structure in IPv6 Addresses. In *ACM Internet Measurement Conference (IMC)*.
- [16] Kensuke Fukuda and John Heidemann. 2018. Who Knocks at the IPv6 Door?: Detecting IPv6 Scanning. In *ACM Internet Measurement Conference (IMC)*.

- [17] Oliver Gasser, Quirin Scheitle, Pawel Foremski, Qasim Lone, Maciej Korczyński, Stephen D. Strowes, Luuk Hendriks, and Georg Carle. 2018. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *ACM Internet Measurement Conference (IMC)*.
- [18] Fernando Gont and Tim Chown. 2016. RFC 7707: Network Reconnaissance in IPv6 Networks. IETF Request for Comments.
- [19] Google. 2020. IPv6. <https://www.google.com/intl/en/ipv6/statistics.html>
- [20] Lili Hervieu. 2019. MAC Address Randomization: How User Privacy Impacts Wi-Fi And Internet Service Providers. <https://www.cablelabs.com/mac-address-randomization-how-user-privacy-impacts-wi-fi-and-internet-service-providers>
- [21] Robert M. Hinden and Stephen E. Deering. 2006. RFC 4291: IP Version 6 Addressing Architecture. IETF Request for Comments.
- [22] Christian Huitema. 2006. RFC 4380: Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs). IETF Request for Comments.
- [23] Elliott Karpilovsky, Alexandre Gerber, Dan Pei, Jennifer Rexford, and Aman Shaikh. 2009. Quantifying the Extent of IPv6 Deployment. In *International Conference on Passive and Active Network Measurement (PAM)*.
- [24] Vector Guo Li, Matthew Dunn, Paul Pearce, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. 2019. Reading the Tea leaves: A Comparative Analysis of Threat Intelligence. In *USENIX Security Symposium*.
- [25] David Malone. 2008. Observations of IPv6 Addresses. In *International Conference on Passive and Active Network Measurement (PAM)*.
- [26] Tomek Mrugalski, Marcin Siodelski, Bernie Volz, Andrew Yourtchenko, Michael C. Richardson, Sheng Jiang, Ted Lemon, and Timothy Winters. 2018. RFC 8415: Dynamic Host Configuration Protocol for IPv6 (DHCPv6). IETF Request for Comments.
- [27] Austin Murdock, Frank Li, Paul Bramsen, Zakir Durumeric, and Vern Paxson. 2017. Target Generation for Internet-wide IPv6 Scanning. In *ACM Internet Measurement Conference (IMC)*.
- [28] Thomas Narten, Richard Draves, and Suresh Krishnan. 2007. RFC 4941: Privacy Extensions for Stateless Address Autoconfiguration in IPv6. IETF Request for Comments.
- [29] Ramakrishna Padmanabhan, Amogh Dhamdhere, Emile Aben, KC Claffy, and Neil Spring. 2016. Reasons Dynamic Addresses Change. In *ACM Internet Measurement Conference (IMC)*.
- [30] David Plonka and Arthur Berger. 2015. Temporal and Spatial Classification of Active IPv6 Addresses. In *ACM Internet Measurement Conference (IMC)*.
- [31] Zhiyun Qian, Z. Morley Mao, Yinglian Xie, and Fang Yu. 2010. On Network-level Clusters for Spam Detection. In *Network and Distributed System Security Symposium (NDSS)*.
- [32] Christian Rossow. 2014. Amplification Hell: Revisiting Network Protocols for DDoS Abus. In *Network and Distributed System Security Symposium (NDSS)*.
- [33] Spamcop. 2020. <https://www.spamcop.net/>
- [34] Spamhaus. 2020. <https://www.spamhaus.org/>
- [35] Kurt Thomas, Rony Amira, Adi Ben-Yoash, Ori Folger, Amir Hardon, Ari Berger, Elie Bursztein, and Michael Bailey. 2016. The Abuse Sharing Economy: Understanding the Limits of Threat Exchanges. In *International Symposium on Research in Attacks, Intrusions and Defenses (RAID)*.
- [36] Shobha Venkataraman, Subhabrata Sen, Oliver Spatscheck, Patrick Haffner, and Dawn Song. 2007. Exploiting Network Structure for Proactive Spam Mitigation. In *USENIX Security Symposium*.
- [37] Yinglian Xie, Fang Yu, Kannan Achan, Eliot Gillum, Moises Goldszmidt, and Ted Wobber. 2007. How Dynamic are IP Addresses? *ACM SIGCOMM Computer Communication Review* 37, 4 (2007), 301–312.
- [38] Dong Yuan, Yuanli Miao, Neil Gong, Zheng Yang, Qi Li, Dawn Song, Qian Wang, and Xiao Liang. 2019. Detecting Fake Accounts in Online Social Networks at the Time of Registrations. In *ACM Conference on Computer and Communications Security (CCS)*.
- [39] Sebastian Zander, Lachlan L. H. Andrew, Grenville Armitage, Geoff Huston, Asia Pacific Network, George Michaelson, and Asia Pacific Network. 2012. Mitigating Sampling Error when Measuring Internet Client IPv6 Capabilities. In *ACM Internet Measurement Conference (IMC)*.

A IMPACT OF THE COVID-19 PANDEMIC ON ANALYSIS RESULTS

A.1 Data Characterization: Prevalence by ASN

Table 1 lists the top ASNs by their percentage of users on IPv6 during a week in April 2020. The relative ranking of top ASNs and their proportions of IPv6 users is similar from January to April, indicating that the pandemic did not heavily impact IPv6 usage for top ASNs. We elide details for other months.

	ASN	Name	Country	IPv6
1	55836	Reliance Jio	IN	96%
2	21928	T-Mobile	US	95%
3	5607	Sky Broadband	GB	95%
4	131445	Adv. Wireless Net.	TH	88%
5	10507	Sprint	US	86%
6	22394	Verizon	US	86%
7	26599	Telefonica Brasil	BR	86%
8	3320	Deutsche Telekom	DE	83%
9	7922	Comcast	US	82%
10	26615	TIM Brasil	BR	82%

Table 1: Top ASNs by their percentage of users on IPv6 during Apr. 13–19.

A.2 Data Characterization: Prevalence by Country

Figure 12 depicts a choropleth of the IPv6 user proportions across different countries for a week in April, for countries with more than 1K users in the user random sample dataset. In Table 2, we list the top countries during a week in January and April. We assess the changes in country-specific IPv6 user proportions over time.

From Jan. 23–29 to Mar. 4–10, we observed minor increases in the IPv6 user proportions for 65 countries, with a median increase of 0.9%. In that same time period, we observed 37 countries with slight dips in their IPv6 user proportions, by a median of -0.5%. For 16 countries, we observed no change. Thus, prior to the lockdowns around the world, IPv6 user proportions shifted only slightly (with a few exceptions discussed below).

Between the week of Mar. 4–10 to Apr. 13–19, we observed 34 countries increase in IPv6 user proportions by a median of 1.2%, while 63 countries dropped by a median of 1.4% (with 13 countries remaining equal). The more prevalent drop in the IPv6 user proportions in many countries was also reflected by our aggregate IPv6 user proportions in Section 4.1.

From Jan. 23–29 to Apr. 13–19, there were some notable country-specific shifts, where IPv6 user proportions change by over 10%.

- Belarus increased in IPv6 proportion by 15.2%, of which 12.2% occurred between Jan. 23–29 and Mar. 4–10. This increase is in line with Belarus’s push for deploying IPv6 country-wide [5] rather than an effect of the pandemic. It is possible though that the pandemic has slowed its rate of increase. Belarus’s IPv6 proportion was growing by about 6% every two weeks through January and February, but this rate slowed to about 1% every two weeks in March and April.
- Germany’s IPv6 proportion increased by 19.4%, from 39.1% to 58.5%, with the majority of the increase occurring in late March. The time period between Mar. 17–23 and Apr. 13–19 accounted for a 14.5% increase (three-quarters of the increase observed). This change coincides with Germany’s lockdown beginning Mar. 22.
- Puerto Rico’s IPv6 proportion dropped 15.5%, from 53.7% to 38.2%, between Mar. 4–10 to Apr. 13–19. There were no notable decreases between January 23–29 and Mar. 4–10 (rather, a 3.3% increase). Thus, this drop may be due to the pandemic.

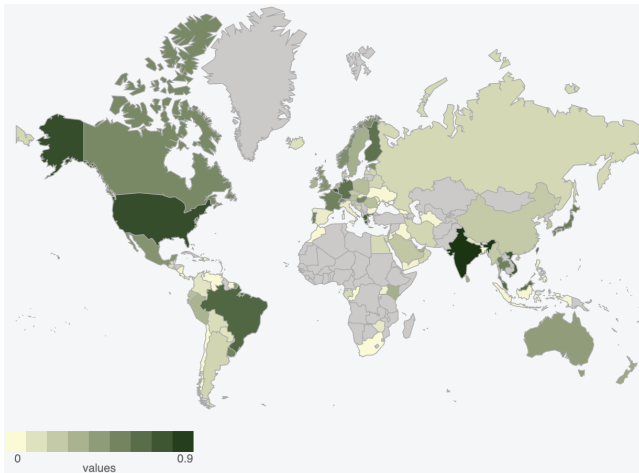


Figure 12: Choropleth depicting the percentage of users seen on IPv6 per country from Apr. 13–19.

Jan. 23–29			Apr. 13–19		
	Country	IPv6		Country	IPv6
1	India	83.4%	India	83.8%	
2	Greece	73.1%	US	73.8%	
3	US	72.2%	Belgium	71.2%	
4	Vietnam	71.2%	Vietnam	70.7%	
5	Belgium	70.2%	Greece	67.8%	
6	Taiwan	68.0%	Taiwan	66.9%	
7	Brazil	66.5%	Brazil	62.9%	
8	Malaysia	63.2%	Malaysia	61.0%	
9	Portugal	55.1%	Germany	58.5%	
10	Finland	55.1%	Finland	53.4%	

Table 2: Top countries by their percentage of users on IPv6 during Jan. 23–29 and Apr. 13–19.

A.3 User-Centric Behavior: IP Addresses per User

We observe that IP diversity among users over time was slightly higher pre-pandemic (i.e., during the early weeks of February) compared to April, for both IPv4 and IPv6. However, the differences are small, and the overall patterns observed in Figure 2 hold. This observation of decreased IP diversity during the pandemic is reasonable as users move around less under lockdowns. However, IP diversity continued to increase over time during the pandemic at a similar rate, for both IPv4 and IPv6, suggesting that much of users’ IP diversity arises from networks assigning new IP addresses on a regular basis, rather than users switching between various networks. For IPv4, this address reassignment is likely due to NATting effects, whereas for IPv6, it is likely due to temporary address assignments by privacy-extended SLAAC [28] and DHCPv6 [26].

A.4 User-Centric Behavior: IPv6 Prefixes per User

As discussed above, IP address diversity was slightly lower during the pandemic. Similarly, we observe an overall small decrease in IPv6 prefix diversity during the pandemic. For prefixes longer than

/64s, there is no significant change, indicating that the higher prefix diversity observed before the pandemic was largely due to more diversity among shorter prefixes. As shorter prefixes are more likely to represent different networks, this suggests that the minor changes in IP diversity during the pandemic were due to users being on fewer networks, plausibly due to reduced mobility.

A.5 User-Centric Behavior: IP Life Spans for Users

We observe that for both IPv4 and IPv6, IP and prefix life spans are slightly longer during the pandemic. However, the differences are minor. For example, comparing Apr. 13–19 data in Figures 5 and 6a to the equivalent curves from Feb. 12–18, no data point differs by more than 4%. The minor change indicates that IP addresses were slightly shorter lived before the pandemic than during, likely due to users being more stationary during the pandemic.

A.6 IP-Centric Analyses

Our IP-centric analysis in Section 6 relied on the IP and IPv6 prefix random sample datasets. Unfortunately, we did not collect this data prior to Apr. 13, and thus cannot compare directly with pre-pandemic conditions. However, given the limited impact of the pandemic on other analysis results, we expect that our overall findings and conclusions should continue to hold.

B IPV6 WEEKEND AND COVID-19 PANDEMIC EFFECTS

We briefly hypothesize on potential causes for the temporal IPv6 effects observed in Section 4.1. Specifically, we observed that user IPv6 prevalence decreases on weekends and during the pandemic, and request IPv6 prevalence increases.

We start by assuming that users are on either residential, mobile, or enterprise networks. During weekends and the pandemic-induced lockdowns, users shift away from enterprise networks to mobile and residential ones, and some users may shift away from mobile to residential ones (e.g., using the home network when staying at home). If enterprise, residential, and mobile networks exhibit IPv6 prevalence in that increasing order, this shift in networks accessed would result in the observed increase in IPv6 request prevalence, as users spend more of their time (and hence a larger proportion of their requests to Facebook) on the more IPv6-adopting residential and mobile networks. Also, as users access fewer types of networks on the weekend and during the pandemic, some users may not access a network supporting IPv6 that they do during a normal (pre-pandemic) workday, resulting in the drop in IPv6 user proportions.

We note that the weekend effect remains to an extent during the pandemic, suggesting that work-related user behavior is also an aspect here, in addition to what types of networks users access. Additionally, it is possible that Facebook’s user population shifted during the pandemic, such as if a large number of users joined that were only IPv4-capable but were not as active as other users.