WHEREIS: IP Address Registration Geo-Consistency

Robert Beverly
San Diego State University
rbeverly@sdsu.edu

Amreesh Phokeer
Internet Society
phokeer@isoc.org

Oliver Gasser
IPinfo
oliver@ipinfo.io

Abstract

The five Regional Internet Registries (RIRs) provide the critical function of IP address resource delegation and registration. The accuracy of registration data directly impacts Internet operation, management, security, and optimization. In addition, the scarcity of IP addresses has brought into focus conflicts between RIR policy and IP registration ownership and use. The tension between a free-market based approach to IP address allocation and use versus policies to promote fairness and regional equity has resulted in litigation and debate that threatens the existence of the RIR system.

We develop WHEREIS, a measurement-based system to geolocate delegated IPv4 and IPv6 prefixes at an RIR-region granularity, and systematically study *where* addresses are physically used *post-allocation* as well as the extent to which registration information is accurate. We define a taxonomy of registration "geo-consistency" that compares a prefix's measured geolocation against the allocating RIR's coverage region and the registered organization's location. While in aggregate over 98% of the prefixes we examine are consistent with our geolocation inferences, there is substantial variation across RIRs. We find that IPv6 registrations are no more consistent than IPv4, suggesting that structural, rather than technical, issues play an important role in allocations. We solicit additional information on inconsistent prefixes from network operators, IP leasing providers, and collaborate with three RIRs to obtain validation. We further show that the inconsistencies we discover manifest in three commercial geolocation databases. Finally, we focus on AFRINIC as a case study of particular importance. By improving the transparency around post-allocation prefix use, we hope to improve applications that use IP registration data and inform ongoing discussions over in-region IP address use and policy.

CCS Concepts

• **Networks → Network measurement**; **Network management**.

Keywords

WHOIS, RIR, IP Registration, IP Geolocation

ACM Reference Format:
Robert Beverly, Amreesh Phokeer, and Oliver Gasser. 2026. WHEREIS: IP Address Registration Geo-Consistency. In *Proceedings of the 2026 ACM Internet Measurement Conference (IMC '26), October 12–16, 2026, Karlsruhe, Germany*. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3777912.3809139

IMC '26, Karlsruhe, Germany

ACM ISBN 979-8-4007-2327-8/2026/10
https://doi.org/10.1145/3777912.3809139

1 Introduction

To ensure the global uniqueness of public IP addresses, the IANA allocates large contiguous blocks of addresses to Regional Internet Registries (RIRs) who then further assign or re-allocate within their respective geographic region.

There are currently five RIRs, each of which serves a different geographical region of the world. Assignments, allocations, and reassignments are governed by RIR policies [3, 5, 6, 36, 46] that are designed to provide uniqueness, efficiency, and accountability for their respective regions. With the advent of IPv4 address exhaustion [56], these policies have become increasingly important as address space is now a valuable commodity [55]. The intent of regional delegation is explicitly codified by the ICP-2 document on the establishment of a new RIR: "The proposed RIR must operate internationally in a large geographical region of approximately continental size" [9].

While the canonical view is that networks located within a particular region obtain IP addresses from the RIR responsible for that region, operational practice is often much more complex. For instance, a company in the Americas might obtain addresses from the European RIR, but use those addresses for a datacenter in Asia. Such "inconsistencies" may be attributable to business, political, abusive, or legal factors, or may simply be due to operational expedience. The tension between a free-market based approach to address allocation versus more needs-based policies to promote fairness and regional equity has become especially contentious with recent litigation threatening the very existence of the RIR system [2, 33]. From an operational perspective, prefixes that are geographically inconsistent with their registration data can impede network monitoring, security, and management.

In this work, we take a rigorous, measurement-based look at the current state of address registration "geo-consistency" across the five RIRs to: i) quantify the extent to which registry information is accurate and can serve operational needs; ii) increase transparency and help the community better understand where addresses are being used; and iii) inform ongoing discussion and current debate over out-of-region (OOR) address use and policy [2, 33, 35, 65, 66]. The primary contributions of this work include:

- **Analysis of IP registration in WHOIS (§3):** Across the five RIRs we find variation (0.2–15%) in OOR organization registrations between RIRs, and fewer OOR IPv6 prefixes as compared to IPv4.
- **Geo-consistency taxonomy (§4):** We create a taxonomy of prefix registration geo-consistency, covering different characteristics such as RIR location, registered organization location, and physical location. We leverage this taxonomy to identify different levels of geo-inconsistencies using Internet measurements.
- **IPv4 and IPv6 prefix geo-audit (§5):** We perform an active measurement campaign and devise a methodology to determine whether prefixes are used in a physical location consistent with

the RIR's region and registered organization's country. While more than 98% of IPv4 and IPv6 prefixes are fully consistent, we identify hundreds of inconsistent prefixes and find substantial variation between RIRs, with AFRINIC exhibiting the highest fraction of inconsistencies in IPv4 and APNIC and LACNIC in IPv6.

- **Analysis of properties for consistent and inconsistent prefixes (§5.3):** To understand possible contributing factors to inconsistent prefixes, we examine prefix age, size, and type, and discover that most inconsistencies in IPv4 are attributable to legacy prefixes, while inconsistent IPv6 prefixes are more likely to be recently updated.
- **Reproducibility and impact (§6)**: We make our code and measurement data publicly available to facilitate reproducibility. Further, we shared our inconsistent prefix inferences with three RIRs who are using our findings to inform their registration database and process.

By improving the transparency around this aspect of post-allocation prefix use, our hope is to improve applications that use IP registration data and inform ongoing discussions over in-region address use and policy—a crucial aspect of a core component of the Internet's future.

2 Background

IP addresses are fundamental to the Internet's operation. The Internet numbers registry system has three primary goals for IP addresses: (1) allocation pool management, (2) hierarchical allocation, and (3) registration accuracy [26]. The IP address hierarchy is rooted in the Internet Assigned Numbers Authority (IANA) which is managed by the Internet Corporation for Assigned Names and Numbers (ICANN) organization. RFC 1366 first proposed to geographically distribute the registry functionality [22]. Today IANA allocates large (e.g., /8's in IPv4 and /12's in IPv6), unique blocks of IP address space to the Regional Internet Registries (RIRs), of which there are currently five: ARIN (North America), RIPE (Europe, the Middle East, and parts of Central Asia), APNIC (Asia and Pacific), LACNIC (Latin America and the Caribbean), or AFRINIC (Africa and the Indian Ocean). RIRs then further allocate IP address blocks to LIRs (Local Internet Registries such as an Internet Service Provider (ISP) or to NIRs (National Internet Registries), which are organizations responsible for the management of IP address allocations at national level (e.g., CNNIC in China). This distribution affords autonomy to the different RIRs to consider region-specific geopolitical policies and constraints, as well as attempting to balance concerns of fairness and regional resource equity.

2.1 Motivation

Policy: Our work is motivated by the ongoing crisis within the AFRINIC region where significant debate and even legal actions—which may threaten the ongoing viability of the RIR itself—following AFRINIC's attempts to enforce in-region address use. AFRINIC has faced prolonged legal challenges since 2020, primarily triggered by disputes over IP address allocations and allegations of governance failures. The most notable conflict involves Cloud Innovation Ltd., which sued AFRINIC after the registry attempted to reclaim large blocks of IP addresses were allegedly not being utilized for the purposes for which they were issued. The legal battles have led to bank account freezes, operational disruptions, and governance paralysis, raising concerns about Internet number resource stability and RIR accountability in the region. [2, 12, 33, 35, 45, 65, 66].

Table 1: NRO Comparative Policy Overview: membership and out-of-region (OOR) policies across RIRs [50]. A ✓ indicates "allowed", while an ✗ indicates "against" policy. OOR polices vary significantly across RIRs.

	Out-of-region			Out-of-region
	Owner	Infrastructure	Routing	usage
AFRINIC	✗	✗	✗	✗[1]
APNIC	✓	✓	✓	✓
ARIN	✗	✗	✗	✗[2]
LACNIC	✗	✗	✓	✗[3]
RIPE	✓	✗	✓	✓

As noted in the ARIN Number Resource Policy Manual: "The primary role of RIRs is to manage and distribute public Internet address space within their respective regions" [6]. However, RIRs must balance efficient and equitable use of IP addresses with the true need for addresses, as well as real-world operational constraints. Further, the five RIRs have *different* policies with respect to out-of-region address use; Table 1 summarizes pertinent policies from NRO's Comparative Policy Overview [50].

AFRINIC has the most restrictive out-of-region usage policies (currently being in the IPv4 exhaustion phase). The account holder must be legally present in the service region and the infrastructure from which the services are originating must be located in the AFRINIC service region. Out-of-region usage is only allowed if the purpose is to bring back connectivity to Africa. ARIN's policy is quite restrictive. Applicants using IP addresses out-of-region should demonstrate a real and substantial connection with the ARIN region [6]. On the other hand, APNIC and RIPE have the most relaxed out-of-region usage policies and they explicitly allows out-of-region use without restriction for its members.

Network Operations: Our work has been performed in coordination with three RIRs: ARIN, RIPE, and AFRINIC. As discussed in §6, we share our data with representatives at these RIRs and obtain informal validation. Our work is further informed by input from the network operator community at NANOG and RIPE meetings. We have received consistent feedback from the operational community that our work is vital to their applications, work, and policies. Indeed, IP registration data is widely used in network monitoring, security, and management. As one individual at an RIR commented to us: "As the country-codes by RIRs are by far the easiest way to group resources by country, your analysis sheds light on to what extent this is accurate."

Database Maintenance: Network operators rely on the WHOIS as a database of record for coordination, business continuity, and security operations. An employee at one RIR confirmed that, while they

[1] Only allowed for connectivity back to Africa.
[2] Must prove real and substantial connection with the ARIN region.
[3] Allowed for Anycast prefixes

```
NetHandle:      NET-104-148-63-0-1
OrgID:          C05266659
Parent:         NET-104-148-0-0-1
NetName:        WEB-OMEGA-DO-BRASIL
NetRange:       104.148.63.0 - 104.148.63.255

OrgID:          C05266659
OrgName:        Web Omega do Brasil
Street:         Rua do Xareu, qd 13, lote 20
City:           Goiania
State/Prov:     GO
Country:        BR
```

Figure 1: Example ARIN prefix registration record. The corresponding organization's country, Brazil, is within the region of a different RIR. WHEREIS seeks to resolve this by geolocating the prefix's true location.

perform pre-delegation checks for policy and region compliance, they have no mechanisms or regular procedures to check for compliance post-allocation—only performing manual investigations upon receiving a complaint. In addition, countless applications, security appliances, and location services rely on WHOIS registration data to provide coarse-grained geolocation information; Section 6.3 examines the impact of geo-inconsistent prefixes on three commercial geolocation databases.

2.2 Terminology

Each RIR maintains registration information, including for example the assigned organization, mailing address, and points of contact, for numbered resources (Figure 1 provides an example). Accurate registration information is important to the operation and management of the Internet—for instance to facilitate timely coordination during a network attack or outage [49]. This registry information is exposed via a public directory service known as "WHOIS" [13].

There are two primary types of IP address registration. The first is an allocation, for blocks of IP addresses which will later be reallocated or reassigned to third parties. The second is an assignment, for blocks that will not be reassigned, or the recipient is an end-user or end-site. Registries such as ARIN require the reassignment of IPv4 prefixes of /29 or more addresses to be registered via a directory services system such as Shared WHOIS (SWIP) within seven days [57, 68].

In this work, we define two regional locations:

- **Out of Region Organization (OOR)**: If the registered prefix owner organization's physical contact address (mailing address) is in a country that is not part of the RIR's region, we say this is an Out of Region Organization (OOR).
- **Out of Region Prefix (ORP)**: When our geolocation of a prefix places the IP addresses outside the registered RIR's region, we call this an Out of Region Prefix (ORP).

As the registered owner of an address prefix may be a multinational company or the customer of an ISP, we take a conservative view of what constitutes an inconsistency by considering both whether a prefix is OOR and/or ORP. This allows us to define different categories of inconsistency (see taxonomy in Section 4.2) and focus on instances where the registrations are most inconsistent.

3 WHOIS Dataset

Before describing our geo-consistency measurement methodology (Section 4), we first detail and examine important characteristics of the prefix registration data in WHOIS.

We use raw "bulk WHOIS" dumps of the IPv4 and IPv6 prefix registration databases as of October 1, 2024 from all five RIRs. Doing so ensures that we obtain the complete data en masse [7]. Note that the RIRs also publish publicly available delegation files that summarize in a standard format core information about the prefixes they manage including status (allocated or assigned), the country code of the organization to which the prefix is delegated, and the date of delegation [52]. As described by Arouna et al., delegation files can be combined with publicly available WHOIS data, for instance to obtain the registration "status" value missing in public ARIN WHOIS or the "maintainer" in public LACNIC WHOIS [8]. We analyze the differences between the bulk WHOIS data from ARIN and the corresponding public extended delegation statistics file from the same date and find e.g., less than 0.1% disagreement in registered country code. Thus, as we apply for, and receive, research access to the complete (non-public) bulk WHOIS data from the RIRs, we henceforth use the bulk WHOIS data in our analysis. Note that we focus on the five RIRs and exclude WHOIS data of National Internet Registries (NIRs). NIRs usually have their own IP addressing and management policy which is outside the scope of this study.

3.1 Macro Characteristics

The bulk WHOIS data is organized in flat text files containing records that consist of key-value pairs, e.g., "NetRange" and "Country" as shown in Figure 1. Each RIR's database has different schemas and idiosyncrasies, these include for example the address prefix representation, different key names, transferred prefixes, and prefixes from other RIRs. We parse each prefix allocation and skip prefixes the RIR does not manage (but still has listed in their database for completeness e.g., annotated with a "not-managed-by" note). We find a number of fine-grained anomalies in WHOIS records, which we deal with diligently (see Appendix B for details).

Some RIRs do not include a country attribute in registration records, but rather use an "OrgID" and include additional records for organizations (see Figure 1). When an OrgID exists, we map the organization's country to the prefix. ARIN, RIPE, and AFRINIC include organization records, while APNIC and LACNIC do not. In total, we map countries for 3,130,262 prefixes from their respective organization records; almost all of these are within ARIN.

Our WHOIS data contains over 8.2M IPv4 and 0.8M IPv6 prefix registrations across the five RIRs, representing a more granular allocation than in the global BGP table [27]. Figure 2 depicts the cumulative fraction of prefixes as a function of the prefix mask across the RIRs. We find nearly 90% of the registered IPv4 prefixes within ARIN, RIPE, and APNIC have a prefix length of /25 or more specific, suggesting active use of SWIP. In contrast, a larger proportion of registered prefixes within AFRINIC and LACNIC are less specific (i.e., smaller network mask). In IPv6, a majority of registrations across all RIRs have a prefix length of /48 or more specific but the distributions of the prefix lengths differ greatly among the RIRs. In

Table 2: WHOIS dataset (Oct 1, 2024), including proportion of prefixes and addresses with Out of Region Owners.

RIR	IPv4				IPv6			
	Prefix (k)	Out-region Owner (k prefix)	Addresses (M /24s)	Out-Region Owner (M /24s)	Prefix (k)	Out-region Owner (k prefix)	Addresses (M /48s)	Out-Region Owner (M /48s)
ARIN	3,104.4	70.8 (2.3%)	5.2	0.1 (2.5%)	308.6	2.3 (0.7%)	13342.2	1.8 (0.0%)
RIPE	3,275.9	23.4 (0.7%)	2.8	0.0 (1.5%)	359.4	3.6 (1.0%)	6993.3	80.8 (1.1%)
APNIC	1,187.3	3.2 (0.3%)	9.4	0.0 (0.1%)	108.2	1.2 (1.1%)	7443.5	25.4 (0.3%)
LACNIC	475.5	1.1 (0.2%)	0.7	0.0 (0.0%)	28.4	0.1 (0.4%)	1082.2	1.4 (0.1%)
AFRINIC	162.8	23.9 (14.7%)	0.5	0.0 (5.7%)	32.3	0.0 (0.1%)	715.6	0.0 (0.0%)
Total:	8,205.9	122.4	18.6	0.2	836.9	7.3	29576.9	109.4

particular, LACNIC has a substantial prevalence of less specific registrations, with 45% of prefixes being less specific than /48, while in the rest of RIRs such prefixes are much less prevalent, never reaching 17% of registrations and, in the case of ARIN, accounting for only 3.7% of registrations.

3.2 Inter-region Registration

Before performing active prefix geolocation in Section 4, we first examine the location information available within the bulk WHOIS data and analyze the prevalence of Out of Region Organizations (OORs).

Prefixes are regularly transferred between RIRs; these transfers are logged and prior work has characterized such published transfers [41]. In contrast, this subsection examines allocations in RIRs that are registered to organizations outside of the RIR's region. For example, a prefix P belongs to an address block allocated to ARIN, but the registered organization within the ARIN database has a physical address in Brazil. In this case, the organization's country belongs to the LACNIC region—a different RIR. Figure 1 provides one such example WHOIS record.

Table 2 provides a macro-level analysis of the number of IPv4 and IPv6 prefixes (in thousands) in each RIR, as well as the fraction of prefixes and addresses (in /24 equivalents) registered to out-of-region (OOR) organizations. Note that the sum of all addresses across the RIRs is larger than the total IPv4 address space—this is due to overlap between prefixes within different RIRs due to "legacy" registrations.

More specifically, the LACNIC and AFRINIC RIRs are newer than the other RIRs. Historically, ARIN managed those resources in Latin America, while RIPE managed IP allocations assigned in the African continent. The so-called Early Registration Transfer (ERX) project was the concerted effort to transfer these early registrations to these registries dedicated to new geographic regions [37, 40, 58].

For example, AFRINIC is responsible for 154.0.0.0/8, however ARIN has a registered allocation of 154.1.0.0/16 (for the company Goldman Sachs). This is a special case called Early Registration Transfers (ERX), where initially the 154/8 was managed by ARIN but got later transferred to AFRINIC after the latter was setup [40].

ARIN has an appreciable fraction (2.3%) of prefixes with out-of-region registered organizations, and dominates in total volume of both out-of-region registered prefixes and addresses. Given the long history of ARIN and the ERX efforts to transfer addresses, this finding is not wholly unsurprising.

AFRINIC is the most recently formed registry (2004). Nearly 15% of AFRINIC prefixes are allocated to out-of-region organizations, suggesting significant apparent involvement of organizations outside of Africa and possibly use of these resources outside of Africa (a hypothesis we test in the next section). Similarly, AFRINIC leads in out-of-region registrations when accounting by total IP addresses (5.7%).

In IPv6 we also find stark differences in OOR between RIRs and differences compared to IPv4. We find APNIC to have the highest share of OOR registrations (1.1%). AFRINIC, with 15% of all IPv4 prefixes OOR, accounts for a meager 0.1% in IPv6. This hints at the use of African IP resources outside the continent for IPv4 only, possibly linked to the IPv4 address shortage [55, 56]. When looking at the share of OOR addresses normalized by /48 prefixes, we see that RIPE has a higher share of OOR addresses (1.1%) compared to OOR prefixes (1.0%), which indicates that this RIR's OOR prefixes are on average less-specific than /48 and thus contain a larger number of addresses. Contrary, all other four RIRs have a lower share of OOR addresses compared to prefixes, showing the use of OOR prefixes more-specific than /48.

3.3 BGP

Our data analysis confirms previous findings that there are more prefixes within the WHOIS as compared to the BGP [27]. While this implies that many of the WHOIS prefixes are subnets within a larger BGP aggregate, this is not always the case. In this subsection, we analyze the degree of alignment (i.e., the amount of address intersection; see Figure 3) between the WHOIS prefixes across the different RIRs with their corresponding BGP announcement(s).

To perform this analysis, we lookup the WHOIS prefixes within a snapshot of the global BGP table as captured on the same date as the bulk WHOIS data by the Routeviews project [60]. In addition to a simple longest-prefix match, we also enumerate the leaves of the best matching node in the prefix trie for those prefixes that are contained within the WHOIS prefix. To see why this enumeration is necessary, consider one WHOIS prefix for "Sony Interactive": `100.42.96.0/20`. While this prefix is not announced within the BGP, there are 10 subnets of this prefix (of size `/23` and `/24`). The WHOIS prefix in this case is a supernet of the BGP announcements, likely due to Sony's CDN making geographic-specific announcements. Our measurements support this inference, as a target within one /23 is clearly within Europe, while a target within a different

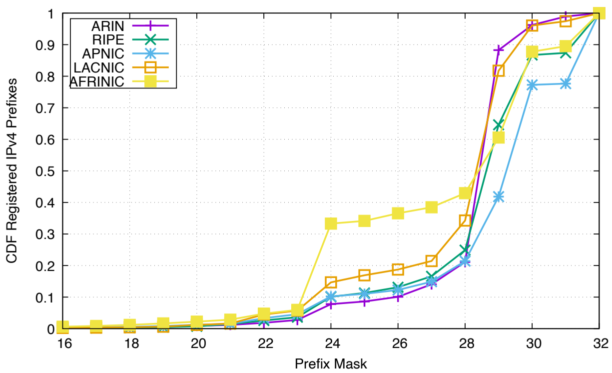

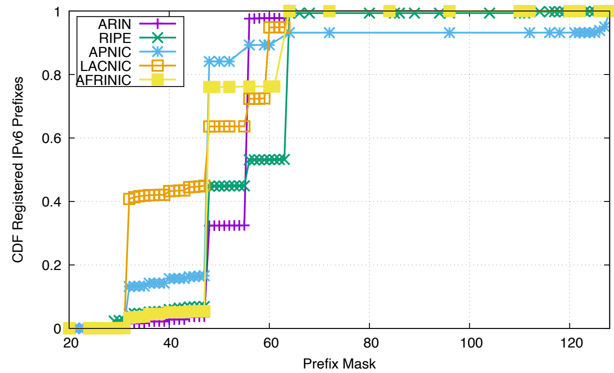

Figure 2: Prefix length distribution of registered IPv4 and IPv6 prefixes per-RIR.

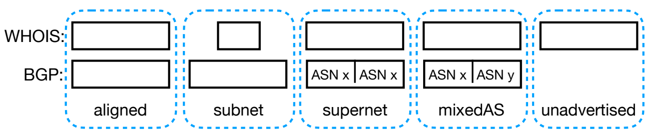

Figure 3: A WHOIS prefix record can align with BGP or be a subnet, supernet, or a supernet containing BGP announcements from different ASNs. In some cases, the WHOIS prefix may not be advertised in the BGP.

/23 of the /20 is clearly within the US. Thus, it is important to differentiate these instances of mis-alignment with the BGP when analyzing our geo-audit results.

Table 7 in Appendix C details the WHOIS to BGP alignment results—the most common case in IPv4, accounting for 95% of WHOIS prefixes, is when the WHOIS prefix is a subnet of a larger BGP announcement. The next most common case in IPv4 are WHOIS prefixes that are the same as the BGP prefix; we term these "aligned." AFRINIC as the largest proportion of aligned prefixes, with nearly 9% of their WHOIS entries exactly matching the BGP announcement. The next case is supernetting, where the WHOIS prefix contains multiple BGP announcements. While only 0.9% of all WHOIS entries are supernets from the same AS, LACNIC and AFRINIC have a higher proportion with 2.9% and 2.4% respectively. Finally, it can be the case that the BGP announcements within the WHOIS supernet are originated by different autonomous systems (ASes)—we term this case "MixedAS"—which occurs for 0.3% of all WHOIS prefixes.

To eliminate cases where inconsistent geolocation for a WHOIS prefix is due to BGP subnetting, we restrict the remainder of our analysis to WHOIS prefixes that are aligned with, or a subnet of, their BGP counterpart.

4 WHEREIS Prefix Geo-Audit Methodology

Having examined prefix registrations as represented within the WHOIS database across the RIRs, we next turn to active prefix geolocation. Our goal is thus to help the community better understand where addresses are being used post-allocation, whether registration information is accurate and can serve operational needs, and

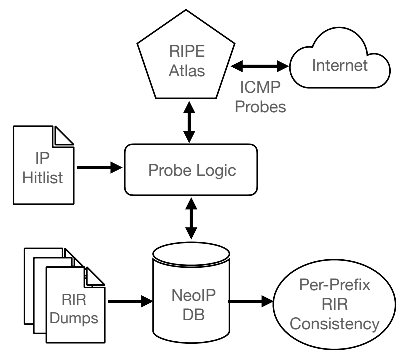

Figure 4: WHEREIS methodology: RIR prefixes are ingested into a database. Probing logic uses IPv4 and IPv6 hitlists to drive RIPE Atlas active probing for constraint-based geolocation and RIR consistency checks.

inform ongoing discussion over RIR policy that are important to maintaining the Internet's open and equitable future.

Figure 4 provides an overview of our WHEREIS geo-audit methodology, with three primary components: (1) ingesting RIR WHOIS into a database; (2) active ICMP probing, which includes selecting the origin vantage point and destination targets within prefixes; and (3) delay-based RIR consistency inferences. Important supporting aspects of this high-level methodology include:

- **Mapping countries to RIRs**: We map countries and their corresponding ISO country codes, to the responsible RIR. Using publicly available information for each registry, we map 244 different countries (including dependent territories with their own ISO code) to the five RIRs using the authoritative NRO mapping [51]. ARIN, RIPE, APNIC, LACNIC, and AFRINIC are each currently responsible for 29, 73, 54, 31, and 57 countries respectively.
- **Target addresses**: Delay-based active IP geolocation requires a responsive address within a prefix. Because it focuses on stable IPv4 addresses, we utilize the ISI Census IPv4 hitlist for this purpose [24]. The hitlist includes a score for each address; to obtain address likely to respond, we filter the hitlist for addresses with a score $\geq$ 99. We then longest-prefix map IPv4 hitlist addresses to their respective RIR prefix to generate candidate targets within fine-grained RIR prefixes for active probing.

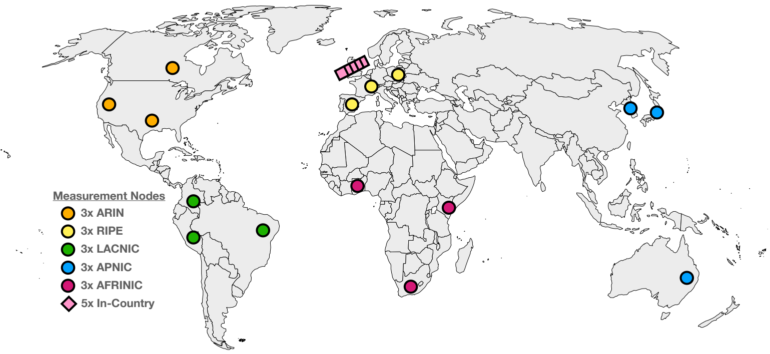

Figure 5: WHEREIS uses RIPE Atlas and constraint-based geolocation to infer true prefix locations. To tighten the feasible geolocation region, we use stable Atlas anchor nodes known to be in each RIR's region as well as nodes within the country of the prefix's registered owner. If the feasible geolocation matches *either* the responsible RIR or the RIR of the owner, we call the registration "consistent," else "inconsistent."

For IPv6, we use ICMPv6 responsive addresses from the IPv6 Hitlist Service [21, 64, 70], specifically data from September 29, 2024, the latest available at the time of our experiments. As aliased prefixes [21]—i.e., a single machine responding to all addresses in a possibly large prefix—are not present in the ICMPv6 responsive addresses, we investigate their prevalence in WHOIS data. We find only 1% of IPv6 WHOIS prefixes covered by the hitlist to be affected by aliased prefixes and thus conclude that using ICMPv6 responsive addresses without aliased prefixes are suitable targets for IPv6.

- **Active Probing**: The accuracy of constraint-based IP geolocation is correlated with the distance from the target to the nearest vantage point; in our case, we desire vantages within the country where the prefix is registered. For this purpose, we utilize the rich coverage afforded by the RIPE Atlas project [11]. As of this writing, Atlas has over 13k "connected probes" (active vantage points) and "anchors" (more stable and maintained nodes) with coverage in over 90% of all world countries [59].
- **Vantage Location Filtering**: While RIPE Atlas nodes include meta-data with their physical location, we follow best practices to ensure that the reported locations are accurate. We first filter out all suspected bad probes in the list published by Izhikevich et al. [30]. Next, for the remaining alive and "connected" nodes, we follow the practice of Darwich et al. [15] and remove probes with bad meta-data, including default country geo-coordinates.
- **Vantage Selection**: From the remaining set of Atlas nodes post-filtering, we select a stable set of 10 vantage points within each RIR region. In addition, we select a stable set of 10 nodes per country—these nodes ensure that we probe the target from *within* the same country as its registered organization. In selecting our vantage points, we choose anchors over probes, and prefer to select those in different ASNs, in so far as possible (for some countries, there are fewer than 10 anchors or probes). As shown in prior literature as well [15, 67], the closest vantage has the most impact in constraining the feasible geolocation space. The choice of 10 within-country vantages and choosing those within different ASNs seeks to maximize the opportunity to find vantages near the target prefix while adhering to our agreement to minimize load on the RIPE Atlas platform, as well as its long-term measurement archival system.
- **Prefix handling**: Central to this work are IP prefixes. We construct a binary radix trie from the registered prefixes and longest prefix match each hitlist target. Thus if a target belongs to two WHOIS prefixes, either in different RIR databases or within the same RIR, we use that target to measure the consistency of the registration information for the most specific prefix—this captures the inherent sub delegation behavior common in prefix registrations. When computing statistics, we take subnetting and prefix aggregation into account so as to not double or over-count.
- **BGP**: We only make inferences for WHOIS prefixes that are either aligned or subnets of the corresponding BGP advertisement (see Table 7) to avoid mischaracterization due to traffic engineering or content distribution.
- **Anycast**: A potential source of error in our geo-audit is anycast [63] wherein a prefix is advertised and reachable from multiple geographic regions. We therefore utilize the current state-of-the-art "AnyCatch" technique and public service [31]. AnyCatch runs its own anycast instance, and issues ICMP probes from different physical locations. If the target is itself not anycast, the ICMP responses should take the single shortest path back to a single instance in the AnyCatch network. However, if the target is anycast, the responses will return to different nodes within AnyCatch. We ignore and do not geo-audit any prefixes that are known to be anycast.

After identifying a likely responsive target within a prefix to geo-audit, we instruct RIPE Atlas to issue ICMP measurements from 20 different vantage points as depicted in Figure 5: three (3) vantages within the region of each of the five (5) RIRs, and five (5) within the country to which the prefix is registered. All of the vantages are selected from the stable and filtered set of vantages described previously. Each measurement is a one-off measurement, public, and tagged to preserve the measurement artifact. We then asynchronously fetch results and insert them into our database.

Result	Description	Example RIR_{Reg}	RIR_{Org}	RIR_{Geo}
(*FC*) Fully Geo-consistent	Geolocates in RIR and org’s region	ARIN	ARIN	ARIN
(*OC*) Org Geo-consistent	Geolocates outside RIR region, within org’s region	RIPE	ARIN	ARIN
(*OI*) Org Geo-inconsistent	Geolocates in RIR’s region, org is OOR	ARIN	RIPE	ARIN
(*RI*) Registry Geo-inconsistent	RIR and org’s region consistent, geolocates OOR	ARIN	ARIN	RIPE
(*FI*) Fully Geo-inconsistent	RIR, org, and geolocation all in different regions	ARIN	RIPE	APNIC

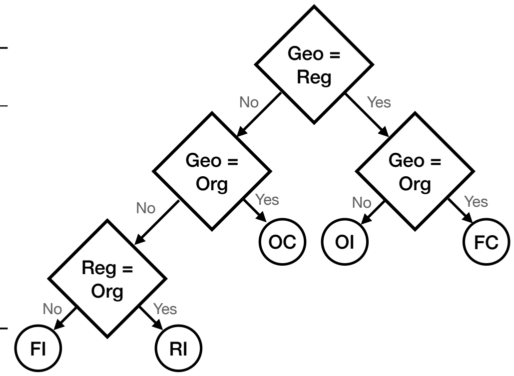

Figure 6: Classification: Given a prefix allocated by RIR_{Reg}, with an organization in RIR_{Org}, we determine the RIR corresponding to the prefixes’ inferred geolocation RIR_{Geo} (§4.1). We then compare the three RIR’s to classify the prefix as either: (FC) Fully Geo-consistent; (OC) Organization Geo-consistent; (OI) Organization Geo-inconsistent; (RI) Registry Geo-inconsistent; or (FI) Fully Geo-inconsistent.

4.1 Speed-of-light-based RIR Inference

As detailed previously in Section 3.2, the organization to which a prefix is allocated may have a mailing address outside of the RIR that manages the prefix’s supernet. For a prefix P allocated to organization Org by RIR_{Reg}, where the contact information for Org lists an address in country CC, we define three important audit values: RIR_{Reg} is the RIR in which the prefix is registered; RIR_{Org} is the RIR responsible for CC, and RIR_{Geo} is the RIR region where P is located based on delay-based geolocation.

From each of the 20 different vantages, RIPE Atlas performs three ICMP echo requests to the target. Among the responses (up to 60 ICMP echo replies) we find the vantage returning the minimum RTT to the target to capture the smallest propagation delay component. As is well-established in the literature, using the minimum RTT is a good proxy for capturing propagation latency (which is fixed) and minimizing effects that contribute variable delay (e.g., queuing and processing delay). We then convert the round-trip latency to a distance based on the 2/3c speed-of-light propagation in fiber optic cables (Appendix D examines the sensitivity of this parameter). From the location of the vantage, we find all physically possible countries and RIRs within that distance radius, i.e., those that do not violate the speed-of-light propagation constraints. Note that while latency-based geolocation can be inaccurate for precision geolocation, conversely it provides a much higher degree of accuracy at the country-level [28]. In this work, we use latency-based geolocation at an even coarser granularity—continents and RIR regions. Further, we conservatively rely on latencies to show that resources are *not* in a particular RIR region.

4.2 Taxonomy of Prefix Registration Geo-Consistency

The value of RIR_{Reg} and RIR_{Org} can be one of the five different RIRs, while RIR_{Geo} (cf. Section 4.1) is a set of possible RIRs we infer based on the 2/3c latency radius. Given these values, we define five different possible inferences, which are summarized with an example in Figure 6. A “Fully Geo-consistent” result is one where the registered RIR, RIR of the organization’s country code, and the RIR of the IP address geolocation determined by RIPE Atlas all match. A “Organization Geo-consistent” result is one where the prefix’s RIR is different than the RIR of the organization’s mailing address, but the inferred geolocation matches the organization’s country code RIR.

The three remaining results show some form of unexpected inconsistency. “Organization Geo-inconsistent” for instance occurs when our inferred RIR geolocation matches the RIR of the prefix, but not the organization’s responsible RIR. Or, both the delegating RIR and the organization’s responsible RIR match, but the inferred geolocation differs from both, a result we term “Registry Geo-inconsistent.” Finally, if all three values are different, we term the result “Fully Geo-inconsistent.”

4.3 Validation

While we obtain and report on result validation obtained from RIRs and providers in Sections 6.2 and 6.5, we additionally validate our measurement methodology. To do so, we utilize the IP addresses and ground-truth geolocations of Speedtest servers [53] within the different RIR covered regions. While these servers are intended for testing bulk transfer capacity, they are well-connected, globally distributed, responsive to our measurements, publicly published [1], and disjoint from the RIPE Atlas probes.

We select 100 Speedtest servers as targets within the registered IP prefix to which they belong. To find server addresses, we query the Speedtest web API from five different network hosts around the world. We obtain 36, 22, 13, 11, and 9 responsive servers with IP addresses in ARIN, RIPE, APNIC, LACNIC, and AFRINIC registered space respectively (91 total; 9 servers did not respond to ICMP requests). For each of these servers, we utilize the published ground-truth geolocations [1]. We then use the entirety of our measurement and inference pipeline as described in this Section on these targets and registered prefixes.

Table 3 displays our complete inferences. Fully Consistent (FC) inferences indicate that our methodology confirms that the prefixes are connected in a region belonging to the RIR where the prefix is registered, and that the contact information is also in the same RIR.

Table 3: Methodology validation against Speedtest server ground-truth (91 servers). Manual investigation of the Organizationally Consistent (OC) and Registry Inconsistent (RI) prefixes confirms 100% accuracy of the methodology for these inferences.

	ARIN	RIPE	APNIC	LACNIC	AFRINIC
(*FC*) Fully Geo-consistent	31 (86.1%)	19 (86.4%)	11 (84.6%)	10 (90.9%)	8 (88.9%)
(*OC*) Org Geo-consistent	1 (2.8%)	3 (13.6%)	2 (15.4%)	1 (9.1%)	0 (0.0%)
(*OI*) Org Geo-inconsistent	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)
(*RI*) Registry Geo-inconsistent	4 (11.1%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	1 (11.1%)
(*FI*) Fully Geo-inconsistent	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)

One ARIN, three RIPE, two APNIC, and one LACNIC prefix are inferred to be Organizationally Consistent (OC), and we manually verify the correctness of these inferences by querying the WHOIS and comparing against the ground-truth geolocation. Finally, four prefixes registered in ARIN and one prefix registered in AFRINIC are Registry Inconsistent (RI), indicating that WHEREIS geolocates the prefix outside of the registered RIR's region. We again manually verify these against the WHOIS and ground-truth locations. In all five RI cases, we find that these are correct inferences, and also that the registered prefixes are subnetted within the BGP via announcements that are more specific than the size of the WHOIS registration. For instance, the AFRINIC RI inference is for a /20 prefix registered in AFRINIC to an organization with an address in Burkina Faso within a /24 BGP announcement. However, WHEREIS confirms, via constraint based latency, that the true location of the target server in that prefix is in the USA. The ground-truth from Speedtest lists the server in Ashburn, VA. Thus, WHEREIS made a correct RI inference.

4.4 Limitations

There are several potential limitations with our methodology. First, we may be unable to find an ICMP-responsive target within a particular prefix. Second, RIPE Atlas may not have any active probe within the country where a prefix is registered, which might result in false negative results. In these cases, we conservatively classify a prefix as geo-consistent, where in reality it is geo-inconsistent. Our results, thus provider a *lower bound* of geo-inconsistent prefixes. Third, the reported location of the Atlas node may be incorrect. However, our use of anchors and rigorous filtering, as described previously, attempts to mitigate any source of error due to Atlas locations.

5 Results

We randomly select 20% of the IPv4 prefixes and all IPv6 prefixes containing at least one target in the hitlist from each of the five RIRs, for geo-audit. When a prefix contains more than one target, we choose two targets to probe. In this fashion, we balance the objective of obtaining a representative sample while adhering to our agreement with RIPE to not overburden the platform as described in Appendix A.

For transparency and reproducibility, our measurement result sets are publicly available from RIPE Atlas and can be found via the `neo-ip-20241018` tag. Each measurement's meta-data encodes the RIR prefix, date, and other relevant information.

Table 4: RIR prefix filtering pipeline.

	IPv4	**IPv6**
Candidate RIR prefixes	94,755	56,838
IP targets probed	112,659	88,546
Responsive targets	82,253	85,884
Responsive prefixes	72,265	56,134
- Anycast	51	118
- NIR	3,669	7,728
- BGP supernet	6,455	5,203
Prefixes after filtering	62,090	43,085
- Conflicting status	6	85
Final Prefixes	62,084	43,000

5.1 Prefix Filtering

This sampling yields 94,755 distinct candidate IPv4 prefixes and 112,659 targets (addresses within the prefix), and 56,838 distinct candidate IPv6 prefixes with 88,546 IPv6 targets. Of these targets, 82,253 IPv4 and 85,884 were responsive to ICMP echo requests (e.g., due to firewall filters or an offline target in the hitlist), corresponding to 72,265 and 56,134 IPv4 and IPv6 responsive prefixes.

To focus on identifying non-ambiguous instances of RIR policy violation, we filter the candidate prefixes as summarized in Table 4. First, as described in §4, we remove 51 and 118 IPv4 and IPv6 anycast prefixes. We then remove 3,669 IPv4 and 7,728 IPv6 prefixes that belong to an National Internet Registry (NIR) as these may have policies independent from the five RIRs. Next, as described in §3.3, we remove 6,455 IPv4 and 5,203 IPv6 prefixes that are either within a BGP supernet or cover mixed ASes. Note that these numbers represent the stepwise filtering operations and the order of operations impacts the number of prefixes removed at each step as some prefixes belong to multiple categories, e.g., a prefix can be both anycast and within an NIR.

We probe targets within our filtered set of prefixes over a 12 day period from October 9–21, 2024 with RIPE Atlas. We find 6 IPv4 and 85 IPv6 prefixes where the two targets within the prefix produce conflicting geo-consistency inferences. We remove these conflicting status prefixes. The remainder of this section presents geo-consistency classification results for the remaining 62,084 IPv4

Table 5: Overview of geo-consistency results.

Result	IPv4 Prefixes	IPv6 Prefixes
Fully Geo-consistent	60,972 (98.2%)	42,019 (97.7%)
Org Geo-consistent	653 (1.1%)	739 (1.7%)
Org Geo-inconsistent	255 (0.4%)	128 (0.3%)
Registry Geo-inconsistent	166 (0.6%)	106 (0.2%)
Fully Geo-inconsistent	38 (0.1%)	8 (0.0%)

and 43,000 prefixes according to the method of Figure 4 and taxonomy of Figure 6.

5.2 Prefix Geo-Consistency

Table 5 presents the aggregate geo-audit results, while Figure 7 shows the per-RIR classification distribution. We observe that, in aggregate and as a relative proportion of the prefixes audited, the registry information is largely consistent: overall 98.2% of IPv4 prefixes are fully geo-consistent. RIPE exhibits the most consistency (over 99.6%). APNIC and LACNIC are also highly consistent, while ARIN is mostly consistent (98.8%). Furthermore, AFRINIC exhibits markedly less geo-consistency with only 92.2% of prefixes being fully consistent. AFRINIC also leads in fully geo-inconsistent prefixes (nearly 1%).

In IPv6, we find overall a higher share of prefixes being consistent, with 98.9% being either fully geo-consistent or organization geo-consistent. Similarly as in IPv4, the majority of inconsistencies are either organization or registry, with low fractions of fully geo-inconsistent prefixes. When investigating consistency levels across RIRs in IPv6, we see quite a number of similarities but also differences compared to IPv4. For similarities, we see that RIPE and APNIC are also highly geo-consistent in IPv6. On the contrary, AFRINIC is highly geo-consistent in IPv6, whereas ARIN exhibits the lowest fraction of geo-consistent IPv6 prefixes.

Delving into the IPv4 inconsistencies (see Figure 8a), we find that the primary contributor (approximately 75%) to ARIN geo-inconsistencies are prefixes that our audit shows to be in the APNIC or RIPE region. Over 50% of the LACNIC geo-inconsistencies are prefixes that are physically in the United States. The geolocation of geo-inconsistent AFRINIC prefixes is dominated by Asia, but also with an appreciable number of prefixes that geolocate to ARIN and RIPE countries. In IPv6 (see Figure 8b), almost one fourth of ARIN geo-inconsistent prefixes is located in Switzerland, with the remainder being found in countries in the LACNIC (e.g., Brazil, Mexico) and RIPE region (e.g., France, the Netherlands, Germany). Similarly as in IPv4, the majority of inconsistent LACNIC prefixes are located in the ARIN region (US and Canada). Finally, RIPE geo-inconsistent IPv6 prefixes are also dominantly in the ARIN region.

5.3 Prefix Characteristics

First, we examine the relationship of a prefix's age, length and allocation type to its inferred consistency. WHOIS records frequently contain timestamps for when the prefix was first registered as well as last updated. Because not all RIRs provide the first registered time, and because we are most interested in the recency of records, we use the last updated field as a measure of the prefix's age. Figures 9a and 9b show the distribution of IPv4 and IPv6 prefix

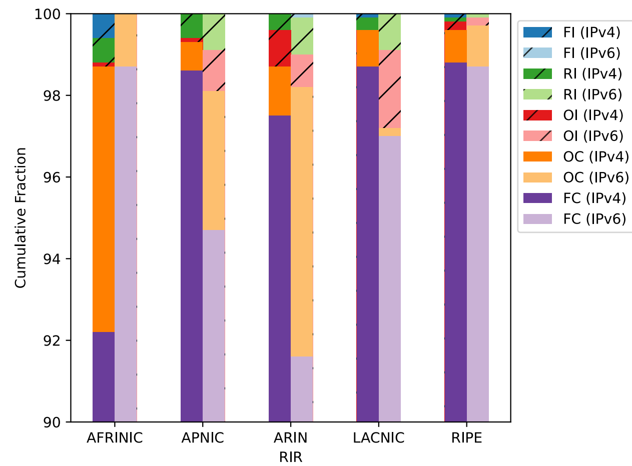

Figure 7: Per-RIR WHEREIS results: >98% of classified prefixes are fully or organization consistent (FC/OC).

last-updated age versus the inferred consistency. In general, and counter-intuitively, prefixes that are more recently updated are more likely to exhibit an inconsistency, although this correlation is more pronounced in IPv6.

Next, we correlate the allocation type to inconsistency results. Figure 10 shows the distribution of allocated, assigned, and legacy (IPv4) or unknown (IPv6) prefixes for each inferred consistency category across IPv4 and IPv6. The majority of fully consistent IPv4 prefixes are assigned, wherein the prefixes are intended for use by the recipient organization. For IPv6, most inconsistencies are attributable to allocated prefixes, or those intended for subsequent distribution by the recipient organization. Legacy prefixes constitute a substantial fraction of fully inconsistent prefixes, suggesting that the registered information for these legacy assignments is not up-to-date. As legacy addresses were allocated by IANA and predate the RIR system, they do not need to adhere to current RIR policies, this is not wholly unsurprising, but still poses challenges to the application and uses of legacy registry data.

6 Impact

In this section we provide details on how geographically inconsistent WHOIS prefixes can impact Internet business models such as IP address leasing, legal proceedings regarding out-of-region transfers for AFRINIC, incorrect geolocation mappings present in popular geolocation databases, the impact of the used measurement platform, and engagement with the wider Internet community.

6.1 IP Address Leasing

Separate from IP address transfers, the scarcity of address space has driven a secondary market for leasing IPv4 addresses. Du et al. [18] recently developed a methodology to infer IPv4 leasing and found 4.1% of all IPv4 address space to be leased. We obtain and use Du et al.'s data to determine the extent to which our inferred geo-inconsistencies are potentially explained by leasing.

Among the prefixes that we infer to be fully inconsistent, we find none of Du et al.'s leased prefixes present for RIPE, LACNIC, or APNIC. 12.5% of the fully inconsistent AFRINIC prefixes and 27.3%

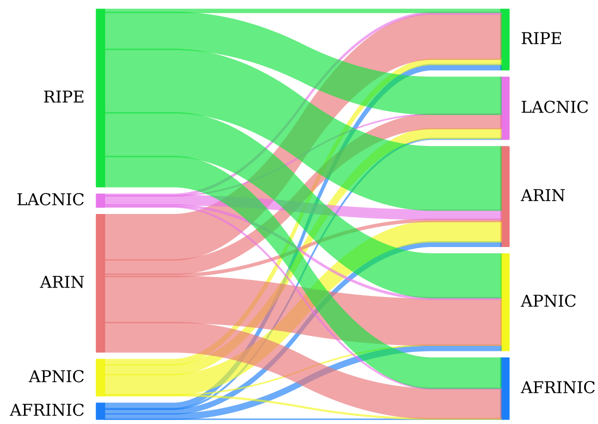

(a) IPv4.

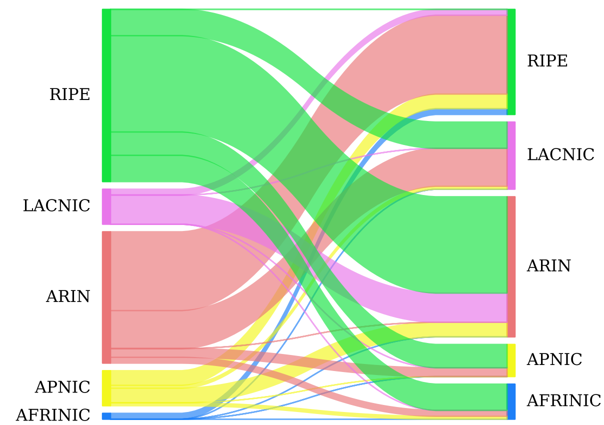

(b) IPv6.

Figure 8: Geo-inconsistent prefixes in IPv4 (left) and IPv6 (right). Source of Sankey is the inconsistent prefix's RIR, Sankey destination is the WHEREIS inferred RIR region.

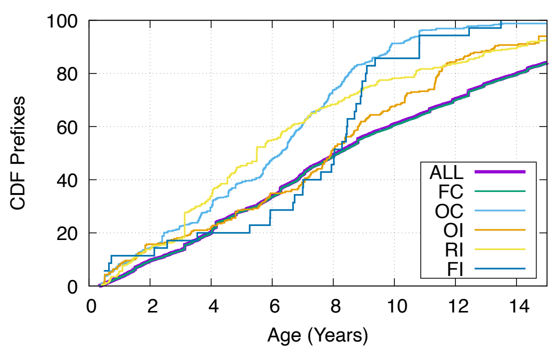

(a) IPv4 last updated

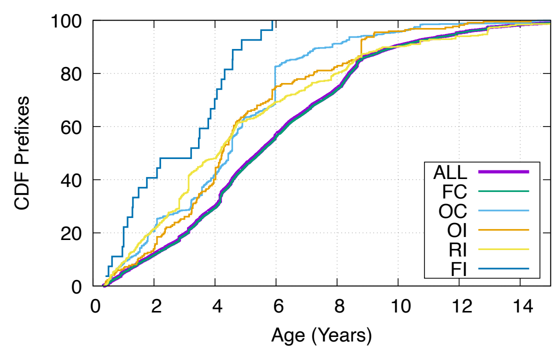

(b) IPv6 last updated

Figure 9: Relationship of geo-consistency to prefix characteristics.

of the fully inconsistent ARIN prefixes are inferred to be leased prefixes. For registry inconsistent prefixes, we find 0.0%, 11.1%, 12.7%, 19.0%, and 42.9% overlap with the leased prefix dataset for LACNIC, APNIC, ARIN, RIPE, and AFRINIC respectively.

These findings suggest that IP leasing is a contributing factor to geo-inconsistent prefixes, but cannot explain the majority of our inferences. We reach out to IPXO, one prominent IP leasing platform provider present in our data and discuss these findings with them (see Section 6.5).

6.2 AFRINIC Case Study

AFRINIC policy stipulates that resources are for the AFRINIC service region, and any use outside the region should be solely in support of connectivity back to the AFRINIC region [3] (see Table 1) . In view of the current litigation [2, 33, 35, 45, 65, 66] and discussions around out-of-region usage in the AFRINIC region, we perform a complete WHEREIS examination of all AFRINIC IPv4 prefixes for which there is a responsive target in our hitlist; in total we are able to classify 7,208 AFRINIC registered prefixes in December, 2024.

Among these prefixes, we find approximately 91% are fully consistent, while approximately 2% are either registry inconsistent or fully inconsistent. Comparing with our randomly sampled measurements in Section 5.2, we find a higher fraction of inconsistent prefixes in AFRINIC. These inconsistent prefixes are part of 51 different prefixes advertised in the global BGP and span 23 different autonomous systems. One BGP prefix accounts for approximately 17% of all inconsistencies, and two ASNs account for more than 53% of the inconsistencies we discover. We omit naming these specific ASN, as one of the stipulations of our interaction with the RIRs was to only report aggregate results and not identify individual ASNs found to be inconsistent.

Due to the timeliness of our work, we share these inconsistencies with representatives of AFRINIC. The two largest contributing ASNs were known to AFRINIC and they validated our findings.

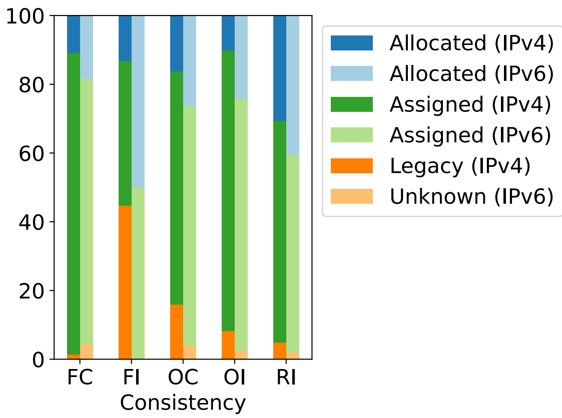

Figure 10: Geo-consistency across prefix allocation type.

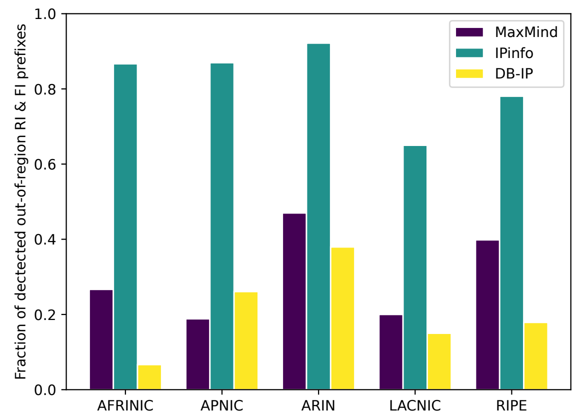

Figure 11: Fraction of correctly detected out-of-region registration inconsistent (RI) and fully inconsistent (FI) prefixes by RIR and geolocation provider. IPinfo has the highest fraction of detected OOR prefixes, followed by MaxMind and DB-IP.

Out of a sample of 106 prefixes shared, 69 were found Fully Geo-Inconsistent (FI) and 37 Registry Inconsistent (RI). AFRINIC staff noted that 19 prefixes were known and of accepted use (Anycast, Global backbone nodes, peering nodes, VSAT, etc), 77 prefixes were known and of unaccepted use as per AFRINIC policies, while 10 prefixes were unknown to them bearing characteristics of both acceptable and non-acceptable use. **Note that we did not specifically target the measurement of any network or prefix, rather these inconsistencies were identified as part of our complete audit of AFRINIC address space—suggesting that WHEREIS correctly identifies geo-inconsistencies**.

6.3 Influence on Geolocation Databases

Next, we evaluate the possible impact of geo-inconsistent RIR prefixes on geolocation databases. To this end, we map registration inconsistent (RI) and fully inconsistent (FI) prefixes using three popular commercial geolocation providers: MaxMind [44], IPinfo [29], and DB-IP [16]. As shown in Figure 11, we find large differences between geolocation providers for identifying out-of-region (OOR) prefixes with detection fractions ranging from 7% to 92% depending on geolocation database provider and RIR. IPinfo has the highest share of detected OOR prefixes, ranging from 65% to 92%, while DB-IP has the lowest share with 7%–38%. Looking at per-RIR detection ratios, we find that IPinfo detects the most OOR prefixes across all five RIRs, with MaxMind being in second place in AFRINIC, ARIN, LACNIC, and RIPE, while DB-IP performs second-best in APNIC. Next, we investigate, in which region each geolocation provider does best: All three geolocation providers have the highest detection ratio for ARIN prefixes with 47% for MaxMind, 92% for IPinfo, and 38% for DB-IP, respectively. On the other side of the spectrum, MaxMind has the lowest percentage of detected OOR prefixes for APNIC (19%), IPinfo for LACNIC (65%), and DB-IP for AFRINIC (7%). This shows that IPinfo is able to detect the large majority of OOR prefixes, while two other prominent geolocation providers might still be relying on location information in WHOIS data, which might not necessarily reflect the actual used location of IP addresses.

6.4 Impact of Measurement Platform

Next, we analyze the impact of the used measurement platform on the detection of geo-inconsistent prefixes. We coordinate with a commercial measurement platform provider and compare our results obtained from RIPE Atlas measurements with results using their measurement platform[4]. We instruct them to use the exact same WHEREIS methodology, with the only difference being to substitute RIPE Atlas with their measurement platform and the selected targets to determine geolocation inconsistent prefixes. As the measurement platform provider runs large-scale measurements towards many more targets compared to our RIPE Atlas measurements, they are able to detect 45.7k registry or fully inconsistent IPv4 and 1.5k IPv6 prefixes, compared to our 204 detected IPv4 and 114 detected IPv6 prefixes.

For prefixes audited by both RIPE Atlas and the measurement platform provider, the commercial provider has better detection recall than Atlas due to the larger scale of measurements and ability to leverage probe vantages closer to the targets. Out of the 82.0k jointly measured IPv4 prefixes, 777 are detected as registration or fully inconsistent. Of these detected IPv4 prefixes, 25.6% are detected by both, 0.3% by RIPE Atlas only, and 74.1% only by the commercial measurement platform. For the 55.5k jointly measured IPv6 prefixes, 827 are classified as inconsistent. In IPv6, the picture looks similar, with 11.5% both detected, 0.7% by RIPE Atlas only, and 87.8% by the commercial provider only.

While the hitlist of responsive targets and limits of the measurement platform may only provide a lower bound of detected OOR prefixes, RIPE Atlas remains a suitable public platform to detect geo-inconsistent WHOIS prefixes and has the benefits of reproducibility and repeatability.

[4]Their measurement network is distributed across more than 800 servers, located in 370 cities and 120 countries.

6.5 Community Coordination

We attempted to solicit validation and additional information from network points of contact as registered in the WHOIS. Surprisingly, there is substantial variation in the fraction of prefixes that have a registered technical, abuse, or administrative point of contact with an email address between RIRs. Table 6 shows that while 99% of the prefixes within ARIN have at least one registered email point of contact, only 41% of RIPE IPv4 prefixes and 26% of RIPE IPv6 prefixes do. Similarly, fewer than 30% of AFRINIC's prefix registrations contain email addresses. Furthermore, in January 2024, AWS announced that they plan to discontinue WHOIS lookups to retrieve email addresses due to their low success rate [10].

Despite this, we make a best effort attempt to contact individual networks. We aggregate results by email point of contact, and find instances where there are both consistent and inconsistent prefixes belonging to the same point of contact. In this way, our hope is to understand cases where administrators choose a particular registry or policy for different portions of their network. We sent our inferences and results to 16 different network points of contact. Unfortunately, at the time of writing we received no responses.

Throughout this work, engaged with the community to share results and incorporate feedback. In particular, we shared our results with three of the five RIRs, ARIN, RIPE, and AFRINIC. Our interaction with AFRINIC is detailed in §6.2.

ARIN's main feedback was that our identified fully inconsistent prefixes were indeed correctly identified by our methodology, and that they primarily represented "hosted services." In particular, ARIN called out an instance of "a hosting provider using ARIN IP addresses for a datacenter in Tokyo, whose end user is located in Estonia." In other instances, especially where ARIN addresses are in use in Mexico, ARIN was not able to add significant amplifying detail, but hypothesized that the service is being delivered in Southern California or Texas, or that resellers are providing service across the border. ARIN stressed that any mismatch between the registry information and the service delivery location is not a violation of their policy.

Furthermore, as we find that a non-negligible share of our identified inconsistent prefixes are leased prefixes (see Section 6.1). We engaged with one of the largest IP leasing platform providers, IPXO. We shared 10 detected inconsistent AFRINIC prefixes with IPXO for validation and feedback. Of the 10 AFRINIC prefixes our methodology determines five being located in the RIPE NCC region and the other five in the APNIC region. Upon checking with their customers, IPXO updated the WHOIS entry of two prefixes to reflect the change in location that we detected. For another two prefixes the customer disagreed with our assessment, even as we provided additional supporting evidence consisting of traceroute data and additional latency measurements to substantiate our inconsistency findings. For the remaining six prefixes IPXO was unable to get confirmation by the customers regarding location changes. This shows that personal contacts and engagement with the wider Internet community, an IP leasing platform provider in this case, can have measurable impact by correcting WHOIS entries, although it remains difficult to do this at scale.

Table 6: Fraction of prefixes with registered email point of contact (POC).

	Prefixes with Email POC (%)	
RIR	IPv4	IPv6
ARIN	99.0	99.2
RIPE	41.4	26.3
APNIC	99.1	98.4
LACNIC	94.9	91.0
AFRINIC	29.1	23.5

7 Related Work

While the present work focuses on IPv4 and IPv6 addresses, the WHOIS protocol is also used to provide access to domain name and autonomous system registration data. Prior work shows that domain names are frequently used for abusive or malicious purposes and registration behavior can be indicative of such misuse [39]. Lauinger et al. analyze the re-registration of domain names after their expiration and show how attackers can leverage residual trust by capturing these expired domains [38]. More recently, Lu et al. dive into WHOIS in the era of GDPR to better understand domain registration privacy [43]. Nemmi et al. examined autonomous system registrations in the WHOIS, and compared registrations against BGP announcements over time [47]. We similarly examine the correspondence between WHOIS and BGP, but at the addressing granularity. Although these prior works all utilize WHOIS registration data, they do not consider IP address block allocations.

Hsu et al. [27] analyzes IPv6 prefixes in different data sources, including WHOIS databases. They find that the most common prefix lengths in WHOIS databases are /64, /48, and /56—with substantial differences between RIRs, with /64 prefixes prevalent in RIPE compared to the dominance of /48 prefixes in APNIC and AFRINIC. However, Hsu et al. do not investigate the location information in the WHOIS databases or the geolocation of prefixes.

Dainotti et al. [14] analyze both passive and active data, including WHOIS registrations and BGP routing, to better understand the utilization of the IPv4 address space. This effort used a commercial geolocation service to locate each /24, and showed the state of IPv4 utilization at the RIR and country levels in 2013. Our work instead explicitly considers current WHOIS georegistration inconsistencies, for both IPv4 and IPv6.

The scarcity of IP addresses has created markets where addresses can be transferred and sold as a commodity [55]. Livadariu et al. analyzed transfers published by the RIRs to characterize the size of prefixes and their eventual use as evidenced in the global BGP table [41, 42]. In contrast, we focus on understanding the true location where addresses are used, and whether these locations are within or outside of the regions for which the corresponding registry is responsible.

More recently, Du et al. design a methodology to examine WHOIS prefix allocations and their corresponding BGP AS origins to infer subnet leasing [18]. While Du et al.'s work does not examine geolocation, we analyze the extent to which the inconsistencies we discover can potentially be explained by leasing in Section 6.1. While unofficial "under-the-table" transfers where the correct location and registry information is not properly updated may explain some of the location inconsistencies we discover, we leave causal analysis to future work.

Many research efforts in the past two decades consider IP address geolocation [17, 19, 20, 23, 32, 34, 54, 61, 62, 67]. These studies mainly focus on the accuracy and techniques for IP geolocation. In contrast, we leverage IP address geolocation to infer inconsistencies of entries in WHOIS databases and their actual physical location. These prior works further highlight challenges in using speed-light constraints for geolocation. While the goal of our paper is not to develop a new or better delay-based geolocation method, we note these prior findings to both analyze the sensitivity to the speed-of-light constraints in our WHEREIS system, as well as guide our section of a conservative inference methodology that prioritizes precision over recall.

More closely related to our work, Zander examined the accuracy of IP allocations against geolocation databases [69] and found a roughly 5% difference by address space. However, this effort, which is over a decade old, assumes the geolocation database as ground truth, as assumption we cannot make as many geolocation databases are themselves constructed in part from WHOIS; we examine the effect of geo-inconsistencies on three commercial geolocation databases in §6.3. To the best of our knowledge, our work is the first to systematically analyze geo-inconsistencies in the WHOIS.

8 Conclusions

This work took the first steps towards understanding the geographic allocation and use of RIR prefixes. We defined a taxonomy to characterize different levels of geo-inconsistencies. Using RIPE Atlas we ran a measurement campaign to infer the geographical region of prefixes and identified 98% as being geo-consistent. We coordinated with three different RIRs with whom we exchanged data for further validation and discussion. We hope these results can meaningfully contribute to the important policy and technical discussions surrounding IP address allocation, as well as highlight areas where RIRs may wish to work with their constituent membership to improve registration accuracy.

Acknowledgments

We thank RIPE for maintaining the Atlas measurement infrastructure. John Curran, Thomas Krenc, Anita Nikolich, Eric Rye, and John Sweeting provided invaluable early feedback. We express our gratitude to AFRINIC staff for validating our findings for the case study and to Ben Du for providing the IP leasing data.

References

[1] 2025. Speedtest Servers. https://www.speedtest.net/speedtest-servers-static.php.
[2] AFRINIC. 2023. AFRINIC Litigation FAQ. https://afrinic.net/litigation-faq.
[3] AFRINIC. 2024. AFRINIC Consolidated Policy Manual. https://afrinic.net/policy/manual/.
[4] Akamai. 2022. Akamai Technologies Completes Acquisition of Linode. https://www.akamai.com/newsroom/press-release/akamai-completes-acquisition-of-linode.
[5] APNIC. 2024. APNIC Number Resource Policy Manual. https://www.apnic.net/community/policy/resources/.
[6] ARIN. 2024. ARIN Number Resource Policy Manual. https://www.arin.net/participate/policy/nrpm/nrpm.pdf/.
[7] ARIN. 2024. Bulk Whois Data. https://www.arin.net/reference/research/bulkwhois/.
[8] Alfred Arouna, Ioana Livadariu, and Mattijs Jonker. 2023. Lowering the Barriers to Working with Public RIR-Level Data. In *Proceedings of the Applied Networking Research Workshop*. 24–26.
[9] ICANN ASO. 2012. ICP-2: Criteria for Establishment of New Regional Internet Registries. https://www.icann.org/resources/pages/new-rirs-criteria-2012-02-25-en.
[10] AWS. 2024. AWS Certificate Manager will discontinue WHOIS lookup for email-validated certificates. https://aws.amazon.com/blogs/security/aws-certificate-manager-will-discontinue-whois-lookup-for-email-validated-certificates/.
[11] Vaibhav Bajpai and Jürgen Schönwälder. 2015. A survey on internet performance measurement platforms and related standardization efforts. *IEEE Communications Surveys & Tutorials* 17, 3 (2015), 1313–1341.
[12] John Curran. 2021. AFRINIC and the Stability of the Internet Number Registry System. https://www.arin.net/blog/2021/08/27/afrinic-and-the-stability-of-the-internet-number-registry-system/.
[13] L. Daigle. 2004. *WHOIS Protocol Specification*. RFC 3912. IETF. http://tools.ietf.org/rfc/rfc3912.txt
[14] Alberto Dainotti, Karyn Benson, Alistair King, Bradley Huffaker, Eduard Glatz, Xenofontas Dimitropoulos, Philipp Richter, Alessandro Finamore, and Alex C Snoeren. 2016. Lost in space: improving inference of IPv4 address space utilization. *IEEE Journal on Selected Areas in Communications* 34, 6 (2016), 1862–1876.
[15] Omar Darwich, Hugo Rimlinger, Milo Dreyfus, Matthieu Gouel, and Kevin Vermeulen. 2023. Replication: Towards a Publicly Available Internet Scale IP Geolocation Dataset. In *Proceedings of the 2023 ACM on Internet Measurement Conference (IMC '23)*. 1–15.
[16] DB-IP. 2024. DB-IP IP Geolocation Database. https://db-ip.com/db/.
[17] Ben Du, Massimo Candela, Bradley Huffaker, Alex C Snoeren, and KC Claffy. 2020. RIPE IPmap active geolocation: Mechanism and performance evaluation. *ACM SIGCOMM Computer Communication Review* 50, 2 (2020), 3–10.
[18] Ben Du, Romain Fontugne, Cecilia Testart, Alex C. Snoeren, and kc claffy. 2024. Sublet Your Subnet: Inferring IP Leasing in the Wild. In *Proceedings of the 2024 ACM on Internet Measurement Conference (IMC '24)*. 328–336.
[19] Patricia Takako Endo and Djamel Fawzi Hadj Sadok. 2010. Whois Based Geolocation: A Strategy to Geolocate Internet Hosts. In *2010 24th IEEE International Conference on Advanced Information Networking and Applications*. IEEE, 408–413.
[20] Alexander Gamero-Garrido, Kicho Yu, Sumukh Vasisht Shankar, Sachin Kumar Singh, Sindhya Balasubramanian, Alexander Wilcox, and David Choffnes. 2025. Empirically measuring data localization in the EU. In *Privacy Enhancing Technologies Symposium (PETS)*.
[21] Oliver Gasser, Quirin Scheitle, Pawel Foremski, Qasim Lone, Maciej Korczynski, Stephen D. Strowes, Luuk Hendriks, and Georg Carle. 2018. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *Proceedings of the 2018 Internet Measurement Conference*. ACM.
[22] E. Gerich. 1992. *Guidelines for Management of IP Address Space*. RFC 1366. IETF. http://tools.ietf.org/rfc/rfc1366.txt
[23] Manaf Gharaibeh, Anant Shah, Bradley Huffaker, Han Zhang, Roya Ensafi, and Christos Papadopoulos. 2017. A look at router geolocation in public and commercial databases. In *Proceedings of the 2017 Internet Measurement Conference*. 463–469.
[24] John Heidemann. 2023. ISI Internet Census IPv4 Hitlist. https://ant.isi.edu/datasets/all.html.
[25] Honesty Net Solutions. 2024. Honesty Net Solutions. http://hns.net.in/.
[26] R. Housley, J. Curran, G. Huston, and D. Conrad. 2013. *The Internet Numbers Registry System*. RFC 7020. IETF. http://tools.ietf.org/rfc/rfc7020.txt
[27] Amanda Hsu, Frank Li, and Paul Pearce. 2023. Fiat Lux: Illuminating IPv6 Apportionment With Different Datasets. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 7, 1 (2023), 1–24.
[28] B Huffaker, M Fomenkov, and k claffy. 2011. *Geocompare: a comparison of public and commercial geolocation databases*. Technical Report. Cooperative Association for Internet Data Analysis (CAIDA).
[29] IPinfo. 2024. IPinfo Geolocation Database. https://ipinfo.io/.
[30] Katherine Izhikevich, Ben Du, Sumanth Rao, Alisha Ukani, and Liz Izhikevich. 2024. Trust, But Verify, Operator-Reported Geolocation. arXiv:2409.19109 [cs.NI] https://arxiv.org/abs/2409.19109
[31] Ben JoJo. 2023. AnyCatch. https://github.com/benjojo/AnyCatch.
[32] Ethan Katz-Bassett, John P John, Arvind Krishnamurthy, David Wetherall, Thomas Anderson, and Yatin Chawathe. 2006. Towards IP geolocation using delay and topology measurements. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*. 71–84.
[33] Eddy Kayihura. 2022. Court Update - Cloud Innovation Ltd vs AFRINIC. https://afrinic.net/20220215-court-update-cloud-innovation-ltd-vs-afrinic-scr-5c-30-21.
[34] Dan Komosny, Miroslav Voznak, and Saeed Ur Rehman. 2017. Location accuracy of commercial IP address geolocation databases. *Information technology and control* 46, 3 (2017), 333–344.
[35] Brian Krebs. 2019. The Great $50M African IP Address Heist. https://krebsonsecurity.com/2019/12/the-great-50m-african-ip-address-heist/.
[36] LACNIC. 2024. LACNIC Policy Manual. https://www.lacnic.net/680/2/lacnic/lacnic-policy-manual.

[37] LACNIC. 2025. ERX - Early Registration Transfer Project. https://www.lacnic.net/2930/2/lacnic/erx---early-registration-transfer-project.
[38] Tobias Lauinger, Kaan Onarlioglu, Abdelberi Chaabane, William Robertson, and Engin Kirda. 2016. WHOIS Lost in Translation: (Mis) Understanding Domain Name Expiration and Re-Registration. In *Proceedings of the 2016 Internet Measurement Conference*. 247–253.
[39] Nektarios Leontiadis and Nicolas Christin. 2014. Empirically measuring WHOIS misuse. In *Computer Security-ESORICS 2014: 19th European Symposium on Research in Computer Security*. 19–36.
[40] Ginny Listman. 2002. Early Registration Transfer (ERX). https://www.arin.net/vault/participate/meetings/reports/ARIN_X/PDF/erx.pdf.
[41] I Livadariu, A Dhamdhere, and A Elmokashfi. 2017. On IPv4 transfer markets: Analyzing reported transfers and inferring transfers in the wild. *Computer Communications Journal* 111 (Oct 2017), 105–119.
[42] I Livadariu, A Elmokashfi, A Dhamdhere, and k claffy. 2013. A First Look at IPv4 Transfer Markets. In *ACM SIGCOMM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*. 7–12.
[43] Chaoyi Lu, Baojun Liu, Yiming Zhang, Zhou Li, Fenglu Zhang, Haixin Duan, Ying Liu, Joann Qiongna Chen, Jinjin Liang, Zaifeng Zhang, et al. 2021. From WHOIS to WHOWAS: A Large-Scale Measurement Study of Domain Registration Privacy under the GDPR.. In *NDSS*.
[44] MaxMind. 2024. MaxMind GeoIP Databases. https://www.maxmind.com/en/geoip-databases.
[45] Milton Mueller. 2021. A Fight Over Crumbs: The AFRINIC crisis. https://www.internetgovernance.org/2021/08/19/a-fight-over-crumbs-the-afrinic-crisis/.
[46] RIPE NCC. 2024. RIPE Policiesl. https://www.ripe.net/publications/docs/ripe-policies/.
[47] Eugenio Nerio Nemmi, Francesco Sassi, Massimo La Morgia, Cecilia Testart, Alessandro Mei, and Alberto Dainotti. 2021. The parallel lives of autonomous systems: ASN allocations vs. BGP. In *Proceedings of the 21st ACM Internet Measurement Conference*. 593–611.
[48] A. Newton and S. Hollenbeck. 2015. *JSON Responses for the Registration Data Access Protocol (RDAP)*. RFC 7483. IETF. http://tools.ietf.org/rfc/rfc7483.txt
[49] Leslie Nobile. 2017. Status and Solutions for Whois Data Accuracy. https://archive.nanog.org/sites/default/files/3_Nobile_Whois_Data_Accuracy.pdf.
[50] NRO. 2024. Comparative Policy Overview. https://www.nro.net/rir-comparative-policy-overview-2024-q4/.
[51] NRO. 2024. List of Country Codes and RIRs ordered by RIR. https://www.nro.net/list-of-country-codes-ordered-by-rir/.
[52] NRO. 2024. RIR Statistics Exchange Format. https://www.apnic.net/about-apnic/corporate-documents/documents/resource-guidelines/rir-statistics-exchange-format/.
[53] Udit Paul, Jiamo Liu, Mengyang Gu, Arpit Gupta, and Elizabeth Belding. 2022. The importance of contextualization of crowdsourced active speed test measurements. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 274–289.
[54] Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. 2011. IP geolocation databases: Unreliable? *ACM SIGCOMM Computer Communication Review* 41, 2 (2011), 53–56.
[55] Lars Prehn, Franziska Lichtblau, and Anja Feldmann. 2020. When wells run dry: the 2020 IPv4 address market. In *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies*. 46–54.
[56] Philipp Richter, Mark Allman, Randy Bush, and Vern Paxson. 2015. A primer on IPv4 scarcity. *ACM SIGCOMM Computer Communication Review* 45, 2 (2015), 21–31.
[57] RIPE. 1998. Shared WHOIS Project (SWIP). https://ftp.ripe.net/ripe/inaddr/arin-templates/swipinstruction.txt.
[58] RIPE. 2005. RIPE NCC to AFRINIC Transition. https://www.ripe.net/community/internet-governance/internet-technical-community/the-rir-system/afrinic/ripe-ncc-to-afrinic-transition/.
[59] RIPE NCC. 2025. RIPE Atlas. https://atlas.ripe.net/coverage/.
[60] Routeviews. 2025. University of Oregon Route Views Project. http://www.routeviews.org/routeviews/.
[61] Quirin Scheitle, Oliver Gasser, Patrick Sattler, and Georg Carle. 2017. HLOC: Hints-based geolocation leveraging multiple measurement frameworks. In *2017 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 1–9.
[62] Yuval Shavitt and Noa Zilberman. 2011. A geolocation databases study. *IEEE Journal on Selected Areas in Communications* 29, 10 (2011), 2044–2056.
[63] Raffaele Sommese, Leandro Bertholdo, Gautam Akiwate, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, KC Claffy, and Anna Sperotto. 2020. Manycast2: Using anycast to measure anycast. In *Proceedings of the ACM Internet Measurement Conference*. 456–463.
[64] Lion Steger, Liming Kuang, Johannes Zirngibl, Georg Carle, and Oliver Gasser. 2023. Target Acquired? Evaluating Target Generation Algorithms for IPv6. In *Proceedings of the Network Traffic Measurement and Analysis Conference (TMA)*.
[65] Jan Vermeulen. 2019. The big South African IP address heist. https://mybroadband.co.za/news/internet/318205-the-big-south-african-ip-address-heist-how-millions-are-made-on-the-grey-market.html.
[66] Jan Vermeulen. 2021. Internet addresses worth R1.8 billion seized. https://mybroadband.co.za/news/internet/405640-internet-addresses-worth-r1-8-billion-seized.html.
[67] Yong Wang, Daniel Burgener, Marcel Flores, Aleksandar Kuzmanovic, and Cheng Huang. 2011. Towards {Street-Level} {Client-Independent} {IP} Geolocation. In *8th USENIX Symposium on Networked Systems Design and Implementation (NSDI 11)*.
[68] C. Weider and R. Wright. 1993. *A Survey of Advanced Usages of X.500*. RFC 1491. IETF. http://tools.ietf.org/rfc/rfc1491.txt
[69] Sebastian Zander. 2012. On the accuracy of IP geolocation based on IP allocation data. (2012).
[70] Johannes Zirngibl, Lion Steger, Patrick Sattler, Oliver Gasser, and Georg Carle. 2022. Rusty Clusters? Dusting an IPv6 Research Foundation. In *Proceedings of the 2022 Internet Measurement Conference*.

A Ethics

We applied for access to ARIN's bulk WHOIS data under their research use provision and abide by ARIN's acceptable use policy. The remaining four WHOIS databases are publicly available. Our work does not involve human subjects or personally identifiable information. Our measurements use RIPE Atlas to send relatively innocuous ICMP probes and follows best practices for active Internet measurements. Furthermore, we contact the RIPE Atlas team to discuss our measurement approach with them. To reduce the load on the RIPE Atlas system, the RIPE NCC suggested that we measure 20% of prefixes for IPv4 and all prefixes in IPv6. Finally, we respect RIPE Atlas resource limits and perform exponential back-off in case the RIPE Atlas API returns an error. As such, this paper does not raise ethical issues.

B Anomalous WHOIS Records

There are several classes of anomalies within the RIR WHOIS records that bear note. First, there exist duplicate prefixes within ARIN; this seems to be most common due to organizational changes, e.g., Akamai acquiring Linode [4]. We find no duplicate prefixes for other RIRs.

Second, the use of "inetnum" and "NetRange," where the record has a starting and ending IP address, permits ranges that cannot be covered within a single contiguous power-of-two prefix. ARIN, RIPE, APNIC, and AFRINIC have 17.8k, 33.1k, 21.2k and 857 of these non-power-of-two WHOIS records respectively.

Third, in the "status" field, which indicates e.g., if a prefix is allocated or assigned, we find variation in the value's textual representation. For example, in APNIC there are nine representations for "Assigned Non-portable" with different variations of capitalization and spacing and even "typo capitalizations" (e.g., "ASSIGNEd NON-PORTABLE"). We surmise that such oddities are a result of manual record entry—a surprising finding given the number of records and the widespread use of the WHOIS, requiring accurate records.

We ignore pathologies in the RIR data, the most common of which are circular references. For example, an RIR may list a prefix as transferred to a different RIR, while that RIR lists the same prefix as belonging to the original RIR. While these types of errors could be due to time differences in the data dumps, the "Updated" timestamps on the records suggest they are simply errors that we cannot resolve.

Finally, there is inter-RIR inconsistency in record fields and values. For instance, RIPE records use "ASSIGNED PA", APNIC uses "ASSIGNED PORTABLE", while LACNIC uses "assigned." Such differences highlight the organic and ad-hoc way in which RIRs have evolved, and significantly complicate not only data processing, but a cohesive understanding of the ecosystem. Initiatives such as the

Table 7: Alignment of WHOIS prefixes within the BGP

WHOIS∈BGP	ARIN		RIPE		APNIC		LACNIC		AFRINIC		ALL	
	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6
Subnet	96.0%	95.5%	95.6%	92.0%	92.6%	27.7%	93.0%	50.5%	84.8%	94.2%	95.0%	84.5%
Aligned	1.9%	1.7%	2.5%	2.7%	3.0%	8.7%	2.6%	18.6%	8.9%	2.1%	2.5%	3.5%
Supernet	0.4%	0.4%	0.6%	0.5%	1.7%	1.4%	2.9%	13.3%	2.4%	0.5%	0.9%	1.0%
MixedAS	0.2%	0.2%	0.2%	0.3%	0.4%	0.4%	0.3%	0.5%	0.3%	0.1%	0.3%	0.3%
Unadvertised	1.5%	2.2%	1.1%	4.5%	2.3%	61.8%	1.2%	17.1%	3.6%	3.1%	1.3%	10.7%

IETF's RDAP [48], which promise to move the WHOIS to a more standards-based format both for the data representation and protocol, will be a welcome change to facilitate future research.

C BGP Alignment

Table 7 shows the fraction of prefixes for each RIR that account for the five BGP-to-WHOIS alignment possibilities (see Figure 3 for a graphical representation of these five cases).

In IPv6, the four classes are similarly prominent as in IPv4, while we see much more diversity between RIRs. In APNIC for example, we only have 27.7% of all prefixes in WHOIS as subnets of BGP, with more than 60% of all IPv6 APNIC WHOIS entries not being found in BGP at all. Upon investigation, we find that a single entity—the Indian company Honesty Net Solutions [25]—is responsible for 79% of all entries in APNIC's WHOIS database missing in BGP. These are all /48 subprefixes of `2401:4800::/32` and have almost all been last updated in WHOIS in May 2012.

D Speed-of-Light Sensitivity

As our method is relies on latency-based constraint based geolocation, a potential source of error are measurement nodes that produce latencies larger than expected given their physical location, e.g., due to circuitous routing, last-mile effects, or non-traditional media such as satellite links. Among the RIPE Atlas nodes used across our measurements, we identified five that reside within Starlink's autonomous system, ASN14593. Recall that each prefix is measured by 20 different Atlas nodes. Among the prefixes measured in our study, our system selected a Starlink node as one of the 20 active measurements for 138 prefixes. We re-ran our analysis to exclude these Starlink Atlas probes/results (while keeping the other 19 probes per prefix) and the overall macro results remain unchanged, i.e., none of the Starlink nodes were ever selected as the minimum-RTT node by which to base the Constraint-Based Geolocation (CBG) constraints.

To examine the sensitivity of our results to the selected speed of light parameter, we re-ran our CBG analysis to compare 2/3c versus 1c. We find that using 1c increases the number of fully consistent prefixes by approximately 0.1%.