IPinfo

WHITE PAPER · 2026

STATE OF ADTECH 2026

# The IP Data Layer

Four opportunities to build better advertising data, with a signal you already get.

# Executive Summary

Advertising technology has spent a decade losing its signals. Cookies survive in name and fail in practice. Mobile device identifiers sit behind consent prompts that most users decline. The bid request itself remains self-declared, with the location, the device, and the app being whatever the sender says they are.

Through all of this, one signal has remained constant with every request, on every platform, in every market, because the internet cannot deliver a packet without it: the IP address.

This paper examines what that signal can carry when it is used effectively. It frames the state of adtech through four opportunities, each supported by production-scale measurement:

**01 · TRAFFIC QUALITY**

Residential proxies can make automated traffic look like ordinary consumer traffic. Measuring that invalid traffic exposure adds visibility beyond traditional hosting, VPN, and proxy checks.

**02 · LOCATION**

Declared bidstream location can be checked against evidence instead of taken on faith.

**03 · CONNECTION CONTEXT**

Further understanding of how many devices and users may share an IP address can help adtech use that data without assuming that one address represents one person.

**04 · PLACE**

Venue-level context can add meaning to location by associating internet activity with physical environments such as hotels, airports, campuses, and stadiums.

These opportunities share a common thread. No single participant sees the entire advertising ecosystem, and difficulties emerge in the gaps between those views. The IP layer is one input every participant can explore to fill those gaps. It adds context to the internet activity each one already sees, whether the goal is cleaner supply, sharper targeting, or more accurate measurement. The questions differ by role, but the layer underneath is the same, and using it well improves the quality of whatever product sits on top.

> **$172B**
>
> projected global ad fraud losses by 2028, per Juniper Research [2]

# An Ecosystem of Partial Views

Money in programmatic advertising flows in one direction: from advertiser through agency and DSP, across an exchange, to a publisher. Information flows differently. A DSP sees bid requests but only for the auctions it enters. An exchange sees its own supply but loses sight of the impression after the auction clears. Verification vendors see the samples they are asked to measure. Publishers see their own traffic and rarely anything upstream of it. Every participant operates on a partial view of the ecosystem, and each of those partial views can be accurate, yet still incomplete.

With no single participant seeing the entire path from request to impression, patterns that become clear across the entire ecosystem can be difficult to recognize from any one position. But better context about the internet activity behind a request can make each of those partial views clearer and more valuable.

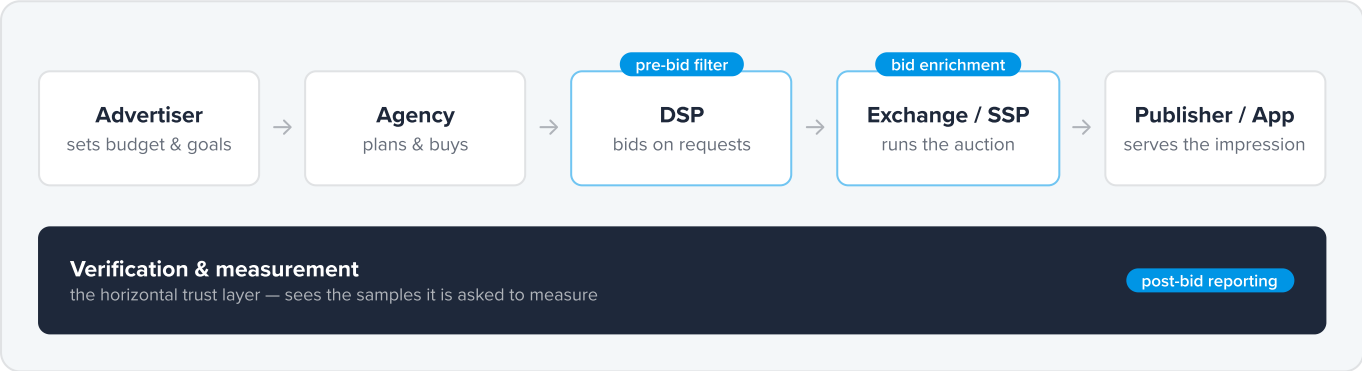


**Figure 1.** The programmatic ecosystem, with markers where IP-level decisions occur. Money flows left to right.

## The Changing Signal Landscape

User identifiers previously helped fill some of those visibility gaps, but that's rapidly changing. Third-party cookies, despite the reversal of formal deprecation, are blocked or degraded across a large share of browsing. Apple's App Tracking Transparency put the mobile advertising identifier behind a prompt, and opt-in rates made the identifier unreliable as a planning basis. Privacy regulation keeps narrowing what may be collected and joined. Meanwhile, the fastest-growing supply, connected TV, arrives largely over IPv6 home connections with no cookie and no per-user identifier at all.

At the same time, the traffic itself is becoming harder to interpret. Automated traffic reached 51% of web traffic in Imperva's 2025 measurement [1], while Juniper Research projects global ad fraud losses of $172 billion by 2028 [2]. There is a widening gap in understanding, with more traffic complexity and fewer consistent signals to interpret it.

> **51%** of web traffic was automated in Imperva's 2025 measurement [1]

## The Signal That Remains

The IP address, however, has remained a constant. It's critical to internet infrastructure, appearing on every request as a requirement of the protocol rather than a policy choice. It's present on every request in all connected devices, on the web, in-app, and CTV. An IP address is a durable signal across an otherwise changing landscape.

An IP can provide context about the network behind a request, where that network operates, whether the connection has been associated with anonymizing infrastructure, how shared the address may be, and even the physical environment it serves.

OPPORTUNITY 01

# Traffic Quality You Can Measure

Residential proxies make automated traffic look like ordinary consumer traffic — and that exposure can now be measured.

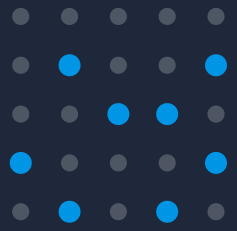

## What Today's Traffic-Quality Checks See

Many established traffic-quality checks are built to recognize familiar infrastructure and behavior like known bots, datacenter ranges, VPNs, and open proxies. Residential proxies are harder to identify because the traffic moves through genuine consumer connections.

A residential proxy network pays real people, or secretly bundles software into free apps, to resell a slice of their home or mobile connection. Someone else's traffic then exits through that connection and reaches the bid stream on a genuine consumer IP (for example, a cable subscriber in Ohio, a mobile customer in Manila). The address geolocates correctly, belongs to a real household, carries no reputation history, and passes every standard check. Fraud operators pay for these networks precisely because the IP looks like a customer.

The Media Rating Council divides invalid traffic into two tiers, and has included residential proxy traffic in the harder tier for years. However, measurement of that traffic has lagged while its usage has exploded in popularity.

**GIVT · GENERAL**

What routine checks catch: declared crawlers, datacenter ranges, known bot signatures.

**SIVT · SOPHISTICATED**

What requires advanced analysis. The MRC places residential proxy traffic in this tier.

## The Measurements

The findings come from IPinfo's adtech benchmark analysis: live programmatic supply measured with IPinfo data at production scale, spanning an unfiltered feed with regional market cuts and supply-chain samples from multiple supply partners, filtered by multiple leading verification vendors. Sources are anonymized and results are reported in directional form. Together they demonstrate the size of the residential proxy visibility gap.

**Residential proxy exposure remains after verification.** Filtering invalid traffic is already standard practice. Most platforms and buyers rely on specialized verification services to remove invalid traffic before it is measured or paid for, and those services handle the well-established categories effectively. Yet across supply samples totaling billions of records, each already filtered by widely used verification vendors, residential proxy exposure appeared in every sample. It ranged from a few percent to more than half of the records in the most exposed sample. Residential proxies also accounted for the large majority of IP-intelligence flags remaining after filtering, while residual hosting and VPN exposure was comparatively small, since those are the categories existing checks were built to recognize [4].

**100%** of verification-filtered supply samples still carried a residential proxy signal [4]

**50%+** of the most exposed sample had cleared verification while carrying the signal [4]

**Residential proxy detection materially changes what becomes visible.** Across billions of unfiltered US bid requests, familiar IP-based categories such as hosting, VPNs, open proxies, and Tor accounted for less than 2% of traffic. Adding residential proxy detection multiplied the share carrying an IP-intelligence flag roughly twentyfold. Separate regional measurements detected residential proxies in more than half of traffic measured across Asia and roughly two-thirds across Africa, approaching nine in ten requests in the most exposed market [3].

**<2%**
of traffic flagged by the standard IP categories alone (hosting, VPNs, open proxies, and Tor)

**~20×**
identified invalid share once residential proxies are detected

**9 in 10**
requests approached in the most exposed market measured

**Exposure is highly concentrated.** At the app level, residential proxy exposure on the unfiltered feed ranged from near zero to nearly all measured traffic. Premium streaming and connected TV inventory was among the cleanest, while several casual gaming and utility apps showed majority exposure, including apps that looked almost entirely clean under standard classification. That variation makes broad averages less useful for individual decisions. The relevant level of exposure depends on the specific supply, market, and traffic being evaluated.

## Why a Flag Is Exposure Rather than a Verdict

An IP address's participation in a residential proxy network reveals something valuable about the connection. It does not mean that every request from that address is invalid. That's because of IP sharing and time of participation.

The internet ran out of IPv4 addresses years ago, and operators responded by pooling many subscribers behind shared public addresses through carrier-grade NAT. How much a network shares depends largely on when it obtained its addresses. Long-established US and European fixed-line providers still assign most homes an address of their own. Meanwhile, mobile operators everywhere, and providers in markets that came online after the address pool emptied, share heavily. The effect is largest across Africa, South and Southeast Asia, and Latin America.

That means a residential proxy observation can affect an IP address that also carries legitimate traffic from other users. The more heavily shared the address, the more important that distinction becomes.

**Time of participation** is also a distinctive factor. Most exit nodes are single-household or mobile connections whose user installed an app that resells idle bandwidth, and the connection acts as an exit node only while that software is active. The same address is one family's ordinary traffic for most of the day and a residential proxy for minutes of it, then rotates out of the pool entirely. A flag observed last week describes an address that may already have left the pool.

That's why residential proxy data is most useful as a flag when it comes with recency and behavior context. **Recency** shows how recently an address was observed in a residential proxy network. **Persistence** shows whether it appears repeatedly or was seen only briefly. Together, those signals support more precise decisions than a residential proxy flag alone.

## Four Seats, Four Responses

The right response to residential proxy exposure depends on where you sit. Four examples:

### EXCHANGE

Screens every incoming bid request at the gate, before the auction, using layered IP intelligence alongside its own supply-quality signals. Because flagging happens at the request level rather than the app level, the invalid request leaves the auction while legitimate inventory from the same application keeps trading.

### AD NETWORK

A quality-positioned network filters aggressively and accepts the short-term volume cost, on the calculation that cleaner delivery raises bid rates, CPMs, and buyer trust over a longer horizon.

### MEDIA-ANALYTICS FIRM

Flags without blocking anything, selling audience-quality reporting to its advertiser clients. Its sharpest internal tell comes before any IP enrichment: campaign traffic that occasionally exceeded the entire addressable market for the product being advertised.

### VERIFICATION PLATFORM

Holds the most direct opportunity of all, since residential proxy coverage extends the SIVT reporting it already provides into the category its customers are currently unable to see.

> **The data remains the same. What changes is the decision each player makes with it.**

OPPORTUNITY 02

# Location You Can Verify

Declared bidstream location, checked against evidence instead of taken on faith.

Location informs decisions throughout programmatic advertising, from campaign eligibility and bidding to compliance and measurement, and much of it enters the ecosystem through the bid request as latitude, longitude, country, and sometimes claimed GPS. Budgets, compliance obligations, and measurement all depend on those declarations, and on the opportunity to verify them.

## How Consensus Erodes Location Data

Verifying a declared location depends on the quality of IP geolocation itself, and a common industry practice quietly works against that quality. Disagreements between sources are expensive: a campaign geofenced to one country gets billed against a provider that places the IP in another, and someone has to absorb the difference. Common practice across adtech is to settle those disputes by consensus, and many providers in turn assemble their data by aggregating one another's answers, so agreement becomes the safest commercial position for everyone involved.

When most of the market averages the same handful of sources, shared mistakes stop looking like mistakes. They harden into the reference answer everyone matches, and a provider that corrects an error looks wrong for disagreeing with the crowd.

## Locating IPs by Evidence, Not Consensus

Evidence changes what a dispute looks like: instead of a vote between opinions, it becomes a question with a measurable answer — a problem the ecosystem can solve rather than settle by averaging.

### How IPinfo ProbeNet locates an IP

Probes ping the address from many cities at once. Each millisecond of round-trip time bounds it to roughly 100 km of that probe; the circles can only intersect where the IP physically is.

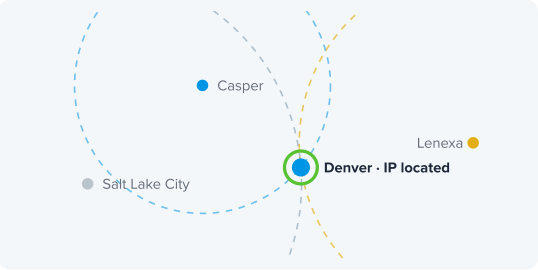


MEASUREMENT AS PROOF

| | |
|---|---|
| **Casper, US** · ping · 589 km bound | 5.771 ms |
| **Salt Lake City, US** · ping · 1,042 km bound | 10.212 ms |
| **Lenexa, US** · ping · 1,144 km bound | 11.213 ms |
| **Denver, US ON-SITE** | 0.070 ms |

Sub-millisecond pings from on-site Denver probes pin the address, with 1,252 further observations as proof.

> **Consensus is a reasonable way to close a billing dispute and a poor way to locate an IP address.**

OPPORTUNITY 03

# Context Without Identity

An IP address describes a building, not the people inside it — and that is exactly what makes it durable.

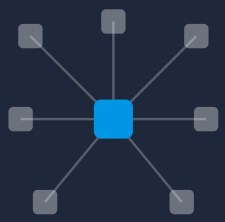

One IP address does not necessarily represent one user. On today's internet, it can represent a household, a mobile gateway, or thousands of devices sharing the same public address.

Research published this year by Stanford's Empirical Security Research Group, using measurement data from a global CDN, shows that IP addresses cannot be identifiers of individual users — and that sharing has increased every year for a decade [5].

**5%**
of client IP addresses carry over half of the world's web requests

**44%**
of addresses show two or more distinct user agents

**<0.2%**
of domains resolve to a dedicated IPv4 address

Stanford Empirical Security Research Group, global CDN measurement, 2026 [5]

## The Purpose of an IP Address

An IP address is the online equivalent of a street address. It is necessarily public; it tells you about the building, and it tells you very little about who is inside. Its job is to tell the network where to deliver something, not who asked for it. Over-blocked households, residential proxy flags read as verdicts, and one-IP-one-user assumptions all descend from treating a delivery address as if it were an identity. The address remains useful. The assumption about what it represents must change.

Every question that matters here is a question about the connection, never about the person using it: is it a datacenter, a carrier gateway, a rented exit node, a venue, and how many devices sit behind it? None of those answers requires a name, an email, a device graph, or a consent chain, and the datasets involved carry none of these. That makes network context useful precisely in the environments where collecting or connecting user-level identifiers is increasingly constrained. The industry does not need to know who someone is to know that a request deserves scoring, that a location claim fails physics, or that an address serves a stadium.

## IP Sharing as a Measurable Signal

Working with an address still demands precision, and sharing is where precision gets lost today. Frequency caps counted per IP overserve every household behind a carrier-grade NAT gateway. Fraud and identity systems that weight IP reputation over-flag mobile gateways, where one address can front thousands of subscribers, and the false positives concentrate in exactly the markets where sharing is heaviest. Household graphs for CTV lean on the IP as their join key at the moment more homes than ever share addresses. Each of these systems is right to use the IP and wrong whenever it assumes one address means one user.

If one IP can represent one device, one household, or thousands of users, there's an opportunity to make the sharing itself measurable. Per-address context can distinguish how many devices currently sit behind an address, if it's a single household line or a carrier gateway, or if it's also acting as a residential proxy exit node. With a device count attached, a frequency-capping system treats a 200-device gateway differently from a family broadband line, a fraud model normalizes risk by the size of the crowd instead of penalizing it, and a household graph weights its join by how shared the address really is. **The assumption becomes a variable.**

An IP address is not a user and should never be read as one. It is context, not identity: the network, the degree of sharing, and the characteristics of the connection add meaning to the traffic each participant already sees, no person required.

OPPORTUNITY 04

# Places You Can Target

Building-level context with no GPS permission, no device identifier, and no individual profile.

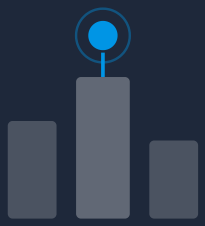

The first three opportunities improve data the industry already uses. The fourth adds a new layer of context. Geolocation describes geographic position through country, region, city, or coordinates. IPinfo Places adds meaning to that position by identifying the physical environment associated with a connection. A connection in Las Vegas can be accurately geolocated to the city. A connection associated with a hotel, airport, convention center, or stadium in Las Vegas carries additional context about where that traffic is coming from. The location sharpens from city level to the building or venue itself.

IPinfo Places maps public Wi-Fi IP addresses to the hotels, airports, stadiums, campuses, and other venues where they originate, with building-level precision. The network is observed directly and tied to a real venue through public venue records, so the tag reflects a place that was verified rather than inferred from an address block. That inverts the usual privacy trade of location targeting: venue-level context requires no GPS permission, no device identifier, and no individual profile, across web, in-app, and connected TV.

**Places as an advertising signal.** Places creates new opportunities while building further trust in geolocation. For targeting, building-level context reaches audiences by verified physical Wi-Fi behavior, like travelers in airports, attendees around a stadium on match day, and guests on hotel networks. For trust, the same data can help evaluate claims about where inventory originates: inventory sold as coming from a premium venue either resolves to that venue's networks or it does not, which turns a marketing claim into a verifiable one.

| 56 | 2.1M+ | 237 |
|---|---|---|
| venue categories | tagged venue IPs | countries |

## THE PLACES TAXONOMY

| | |
|---|---|
| TRANSIT & TRAVEL | In-Flight, In Transit, Airport, Travel Agency, Bus Station, Car Rental, Train Station, Subway Station, Ferry Terminal |
| ENTERTAINMENT & NIGHTLIFE | Casino, Music Venue, Theater, Nightclub, Cinema, Entertainment Venue, Event Venue |
| ATTRACTIONS & CULTURE | Museum, Library, Attraction, Zoo, Botanical Garden, Aquarium |
| RETAIL & SHOPPING | Retail, Grocery Store, Home Improvement, Convenience Store, Shopping Mall, Pharmacy, Market |
| PROFESSIONAL & FINANCIAL | Business Services, Home Services, Bank, Real Estate, Laundromat, Self Storage |
| HEALTH & WELLNESS | Personal Care, Veterinary |
| ACCOMMODATION | Hotel, Campground, Mobile Home Park |
| SPORTS & RECREATION | Gym, Sports Center, Golf Course, Stadium, Water Park, Theme Park, Ski Resort |
| FOOD & DRINK | Restaurant, Bar, Coffee Shop |
| BUSINESS & COMMUNITY | Public Space, Community Center, Conference Center, Coworking Space |
| AUTOMOTIVE | Car Services, Car Dealer |

Growing as new venues are measured. Categories describing people at schools, places of worship, or medical facilities are deliberately withheld.

# One Data Layer, Many Roles

Adtech is an ecosystem in which every participant sees a different part of the advertising supply chain. Traffic quality, location, network context, and place add more critical information to those existing views. The IP layer they share gives each participant more context about the internet activity it already sees. The questions and decisions differ depending on their role.

| PLAYER | HOW IP CONTEXT CAN BE USED |
|---|---|
| **DSP / buyer** | Pre-bid filtering on hosting and residential proxy exposure; independent check on declared bidstream location before budget is spent; per-supply-partner and per-market thresholds built on residential proxy share, hosting exposure, and device-count context; venue-level Places context for validating premium-venue inventory claims. |
| **Exchange / SSP** | Screening at the gate, before the auction, at the request level rather than the publisher level: residential proxy flags weighed together with the platform's own supply-quality signals, so requests judged invalid on the combined evidence drop out while legitimate inventory from the same app keeps trading. |
| **Ad network** | Request-level screening on residential proxy and hosting flags as a product characteristic: aggressive filtering traded against short-term volume for higher bid rates, CPMs, and buyer trust over time. |
| **Verification / measurement** | Residential proxy coverage as an extension of existing SIVT reporting, closing the category the MRC defines but the market rarely measures. |
| **DMP / data platform** | Joining IP context — venue and place category, verified geolocation, privacy, and network flags — onto billions of bid requests as an enrichment layer, so a city-level impression becomes a venue- and event-level one for DSPs, SSPs, and other DMPs downstream. |
| **Agency / brand** | Audience-quality reporting on campaigns without blocking anything: how much of the delivered audience shows machine behavior, and where. |
| **Publisher / app developer** | Request-level attribution of invalid traffic, so a developer with a good app keeps earning while the residential proxy-routed requests inside it are removed. |

**Figure 2.** How IP context can be used, by seat in the ecosystem.

The value of the IP layer comes from how it's measured and the context it adds. Observed residential proxy signals reveal more about how traffic reaches the ecosystem. Actively measured geolocation provides an independent view of where that traffic originates. Per-address connection context helps explain what sits behind a shared address. Places adds information about the building-level environment associated with it.

Together, these signals make the internet traffic already visible across adtech easier to interpret. A buyer, exchange, verification provider, data platform, and publisher each approaches that activity with different questions because each has a different role to play.

The practical starting point is to measure your own traffic. Take a representative sample, enrich it with IP context, and break it out by market, supply partner, and app, or by whatever dimensions drive your decisions. Industry-wide patterns like the ones in this paper show what exists at scale; your own traffic shows which signals actually change the decisions you make.

As internet activity grows more complex and familiar signals become less reliable, understanding the traffic itself becomes more important. The IP layer provides a durable foundation for that understanding, adding context to the connections already moving through every part of the advertising ecosystem.

# About IPinfo

IPinfo is the internet data company. The internet has become harder to understand than at any point in its history, with more networks, more layers, more automation, more actors, and no single participant seeing the entire system holistically. IPinfo exists to close that understanding gap. We measure, analyze, and contextualize the internet continuously, turning raw signals into reliable knowledge of how networks, infrastructure, and traffic behave in the real world.

The difference is method. Most IP data is assembled by aggregating registries, geofeeds, and other providers' answers. Aggregation alone can only measure agreement between sources; when the sources share assumptions, consensus reinforces errors instead of revealing them.

IPinfo takes a different path. Our engine ingests signals from hundreds of independent sources and then validates them against physical evidence. ProbeNet, our internet measurement platform of roughly 1,300 servers worldwide, actively probes the internet every day, and the physics is unforgiving, since each millisecond of round-trip time bounds an address to roughly a hundred kilometers [6].

Agreement with other datasets is never the goal. Alignment with internet reality is, and because the method is evidence-driven, our conclusions can be checked rather than merely accepted: a public version of the measurement platform lets anyone test an address for themselves.

The result is a shared layer of context about IP addresses and the networks behind them, including geolocation, ASN and company data, privacy and hosting signals, residential proxy detection, connection characteristics, and Places.

| OPPORTUNITY | IPINFO DATA BEHIND IT |
|---|---|
| **Traffic quality you can measure** | Privacy Detection (hosting, VPN, proxy, relay classification) and Residential Proxy Detection: exit nodes observed directly in proxy networks, refreshed daily, with the operating service named and each address's persistence in the pool tracked over time. |
| **Location you can verify** | IP Geolocation built on evidence: administrative records and device observation, filtered by ProbeNet internet measurement, delivered with a confidence radius on every answer so downstream systems know how far the evidence reaches. |
| **Context without identity** | Per-address network context: carrier and network type, mobile and satellite classification, connection type, and device-count signals that make the crowd behind an address measurable without touching identity. |
| **Places you can target** | Places: venue-level mapping of the addresses serving hotels, stadiums, airports, campuses and public Wi-Fi deployments. |

Every dataset is produced by the same engine and the same evidence standard, delivered at bidding speed through API, downloads, and the major cloud data platforms. The most practical way to see that data in action is to run it against your own traffic: send us a representative sample and we will enrich it and segment it by market, supply partner, app, or the dimensions you care about — showing you where residential proxy exposure concentrates, where declared and independently derived location diverge, how connection characteristics vary across your supply, and where Places context is available. Visit **ipinfo.io/adtech** to get started.

# Annex: Method and Definitions

## Terms

**GIVT and SIVT** follow the Media Rating Council's Invalid Traffic Detection and Filtration Guidelines. General invalid traffic is identifiable through routine list-based and parameter checks; sophisticated invalid traffic requires advanced analytics or human intervention, and includes invalid proxy traffic. This paper uses "fraud" only for intentional operations, never as a synonym for invalid traffic.

**Residential proxy**, as measured here: an IP address classified as clean or mobile by standard privacy detection that is concurrently present in IPinfo's residential proxy database, which is built by direct observation of proxy networks and refreshed daily. The flag denotes exposure of the address, and does not by itself classify any individual request as invalid.

## References


1. Imperva, 2025 Bad Bot Report. Automated traffic accounted for 51% of all web traffic in 2024; malicious bots 37%. imperva.com/resources/resource-library/reports/2025-bad-bot-report
2. Juniper Research, Quantifying the Cost of Ad Fraud 2023–2028, 2023. Losses estimated at $84.2 billion in 2023, projected to reach $172.3 billion by 2028.
3. IPinfo AdTech Benchmark Analysis, 2026: unfiltered-feed evaluation, US feed plus separate regional cuts.
4. IPinfo AdTech Benchmark Analysis, 2026: post-verification supply samples, each filtered by a widely deployed verification vendor before measurement.
5. R. Habib, S. Singanamalla, M. Fayed and Z. Durumeric, On IP Addresses as Identifiers of Internet Users and Services, Stanford Empirical Security Research Group, 2026.
6. IPinfo Geolocation Methodology, ipinfo.io. ProbeNet deployment figures as of August 2026.

IPinfo

NEXT STEP

# Measure your own traffic

Send us a representative sample and we'll enrich every IP with our data: residential proxy, VPN and hosting flags, derived location, the ASN and company behind the IP, and the venue where we've mapped one, across 50+ categories including airports, hotels, transit and retail.

TALK TO SALES

**sales@ipinfo.io**

SEND A TRAFFIC SAMPLE

**ipinfo.io/adtech**