IP Geospatial Intelligence

Exposing the Limitations of IP Geolocation Through Real-World Spatial Observation

Mar 19, 2026

Technical White Paper | Ed Gibbs, Vice President of Research, WHOIS

Executive Summary

IP Geolocation, which is the widely deployed practice of inferring a device's physical location from its assigned IP address is fundamentally predicated on the accuracy of Regional Internet Registry (RIR) registration data and the operational assumptions of static IP-to-geography mappings. This paper presents empirical evidence that these assumptions are structurally invalid in a growing class of deployment scenarios.

Using the IP Spatial platform, which derives device location from observed network telemetry rather than registration metadata, this paper documents two high-confidence case studies in which the **ground-truth physical location** of a device diverges from its RIR-registered country by thousands of kilometers. Across eleven major IP geolocation providers, none detected these discrepancies. The implications for fraud detection, compliance enforcement, digital rights management, and cybersecurity threat attribution are significant.

Background: How IP Geolocation Works

Traditional IP geolocation systems construct their databases from a combination of the following authoritative and heuristic sources:

- RIR allocation records — ARIN, RIPE NCC, LACNIC, APNIC, and AFRINIC maintain IP address block assignments to organizations and network operators, each of which self-reports a geographic jurisdiction at the time of allocation. This reported location is largely static and is not validated against physical infrastructure.
- BGP routing tables — Border Gateway Protocol announcements convey Autonomous System (AS) path information. Geolocation databases cross-reference AS numbers against known operator countries to infer location at the AS or prefix level.
- Active and passive network probing — Round-trip time (RTT) measurements, traceroute path analysis, and reverse DNS lookups are used to refine city- or region-level estimates within a country boundary already established via RIR data.
- Commercial data feeds — Some providers supplement the above with wi-fi hotspot databases, user self-reported location signals, and mobile carrier handoff data.

The fundamental weakness shared by all of these methods is that they are inferential: they reason about where an IP address should be based on administrative and routing metadata, not where the device using that address actually is. IP Spatial, by contrast, derives location from direct spatial observation of the device — specifically, GPS-correlated telemetry collected at the endpoint or from network infrastructure with physical location anchors.

Case Study 1: Colombia-Registered IP Address Observed in Indonesia

Subject IP Address and Registration Metadata

IP Address	200.69.74.28
IP Range	200.69.74.0 – 200.69.74.255
RIR	LACNIC (Latin America and Caribbean Network Information Centre)
Registered ASN	AS262186
Registered Operator	TV AZTECA SUCURSAL COLOMBIA
Operator Hostname	azteca-comunicaciones.com
RIR-Registered Country	Colombia (CO)
IP Spatial Observed Country	Indonesia (ID)
IP Spatial Coordinates	-6.205xxxxxx, 106.692xxxxxx (Nerogtog, Banten)
Geodesic Distance (CO to ID)	~12,399 km

Confidence Level	**HIGH**
Likely Explanation	VPN tunneling or proxy routing via Colombian IP infrastructure

IP Geolocation Provider Comparison

The following table documents the geolocation output from eleven commercially deployed IP geolocation providers queried against IP address 200.69.74.28. All eleven providers placed the address in Colombia, consistent with LACNIC registration data. None detected that the device was physically present in Indonesia.

Provider	Reported Country	Reported City/Region	Reported Coordinates
WHOISXMLAPI	Colombia	Socotá, Boyacá	6.04028, -72.63509
IPinfo.io	Colombia	Soatá, Boyacá	6.3337, -72.6828
IP2Location (DB6)	Colombia	Socota, Boyaca	6.0403, -72.6351
ipinfo.io (via IPLocation.Net)	Colombia	Soatá, Boyacá	6.3337, -72.6828
DB-IP	Colombia	Santa Rosa de Viterbo, Boyacá	5.874, -72.9822
IPregistry.co	Colombia	Tipacoque, Boyacá	6.334, -72.683
IPGeolocation.io	Colombia	Bogotá, D.C.	4.711, -74.072
IPapi.co	Colombia	Santa Rosa de Viterbo	5.874, -72.9822
criminalip.io	Colombia	Duitama, Boyacá	5.824, -73.040
ipapi.is	Colombia	Socotá, Boyacá	6.040, -72.635
iplocate.io	Colombia	Bogotá, D.C.	4.610, -74.082
IP Spatial (Ground Truth)	Indonesia	Nerogtog, Banten, Indonesia	-6.20xxx, 106.69xxx

The discrepancy between the RIR-registered location (Colombia) and the IP Spatial-observed location (Nerogtog, Banten, Indonesia) represents a geodesic distance of approximately 12,399 km. This magnitude of error is not a matter of city-level precision — it represents a continental-scale failure of the underlying geolocation methodology.

Raw Telemetry Record

The following CSV record is the source data used by the IP Spatial platform to derive the ground-truth location for this IP address:

```
v1_range,lat1,lon1,city1,state1,country1,ipgv2,lat2,lon2,city2,state2,country2,
rir_country2,diffDistance,match_type
200.69.74.0-200.69.74.255,6.040xx,-72.635xx,Socotá,Departamento de
Boyacá,CO,200.69.74.28,-6.205xxxxx,106.692xxxx,Nerogtog,Banten,ID,CO,12399,rir_
country_match
```

The match_type field value of “rir_country_match” indicates that the RIR-registered country (CO) and the geolocation-reported country are consistent with one another, while the IP Spatial-observed country (ID) diverges. This is the hallmark signature of a tunneling or proxy scenario in which the traffic origin is obscured at the IP layer. In the above example, the longitude and latitude have been altered to not reveal the exact location.

Case Study 2: Chile-Registered IP Address Observed in China

Overview

A second high-confidence example further corroborates the systemic nature of IP geolocation failure in tunneling scenarios. IP address 186.79.66.140, registered to an operator in Chile, was observed by the IP Spatial platform to be physically located in China across six consecutive weekly observation windows in November 2025. The degree of confidence is 100%, categorized as HIGH.

All IP geolocation providers queried for this address returned Chile as the country. No provider identified the actual physical location as being in China. The coordinates consistently observed by the IP Spatial platform — latitude 34.37xxxx, longitude 109.25xxxx — place the device in Shaanxi Province, China.

Observed Telemetry — Six-Week Window

The IP Spatial platform recorded the following observations for IP address 186.79.66.140. The consistency of the coordinates across all six observation dates eliminates noise and ephemeral routing artifacts as explanatory factors, strongly indicating a persistent physical installation.

Observation Date	Reported Country	Latitude	Longitude
2025-11-25	China (CN)	34.37xxxx	109.25xxxx
2025-11-26	China (CN)	34.37xxxx	109.25xxxx
2025-11-27	China (CN)	34.37xxxx	109.25xxxx
2025-11-28	China (CN)	34.37xxxx	109.25xxxx
2025-11-29	China (CN)	34.37xxxx	109.25xxxx
2025-11-30	China (CN)	34.37xxxx	109.25xxxx

* Longitude and latitude have been altered for privacy

The stability of these readings — identical latitude and longitude across 569, 687, 350, 626, 446, and 112 independent readings respectively indicates a stationary device. This pattern is inconsistent with a transient mobile connection and more consistent with a fixed infrastructure endpoint, such as a server, router, or set-top device, tunneling traffic through a Chile-registered IP range.

Why Traditional IP Geolocation Fails in These Scenarios

The VPN and Proxy Tunneling Problem

When a device establishes a VPN tunnel or routes traffic through a proxy, the IP address visible to the network is that of the tunnel endpoint, not the device's own interface. The tunnel endpoint's IP address is registered to an operator in a specific jurisdiction, and all downstream geolocation inference follows from that registration. The physical location of the originating device is entirely transparent to geolocation systems that lack direct spatial observation capability.

In both case studies above, the IP address being queried belongs to a legitimate operator in the registered country. There is no spoofing, no BGP hijacking, and no WHOIS fraud. The geolocation systems are working precisely as designed and that is precisely the problem. The design assumption that IP registration geography reflects device geography is violated by the existence of tunneling infrastructure.

The RIR Registration Problem

RIR registrations are filed at the time of IP block allocation and are rarely updated to reflect operational changes. An operator who deploys infrastructure internationally does not amend their LACNIC or RIPE NCC records to reflect the physical location of each device using an address within their allocated range. The registration records the administrative owner, not the physical location of traffic origination.

Furthermore, RIR data is self-reported. There is no verification mechanism that audits whether an operator's declared country of operation matches the geographic distribution of their active infrastructure. This is a structural property of the internet's administrative architecture, not a correctable data quality issue.

The Static Database Problem

Geolocation databases are periodically updated, but they operate as retrospective snapshots. Network operators routinely lease, sell, reassign, and tunnel IP ranges across jurisdictional boundaries faster than any database update cycle can track. Even providers with API-based real-time lookups are querying static databases that were compiled from historical RIR and BGP data; they are not performing live spatial observation of the queried device.

IP Spatial Methodology

IP Spatial addresses these limitations through a fundamentally different approach: direct spatial observation of IP addresses in the wild. Rather than inferring location from registration metadata, the platform collects GPS-correlated observations of IP addresses as they appear in real network traffic, aggregated across a large-scale distributed sensor network.

Key characteristics of the IP Spatial methodology include:

- Empirical observation, not inference — Location data is derived from direct measurement of where packets originate, not from administrative metadata about where operators are registered.
- Temporal coverage — Observations are timestamped and accumulated over time, enabling the platform to distinguish transient routing events from persistent physical deployments.
- Reading volume as confidence signal — Each observation record includes a reading count. High reading counts over multiple observation windows (as in the China case study above) increase confidence that the observed location reflects a stable physical installation rather than an anomalous transient path.
- Geodesic discrepancy scoring — The platform computes the geodesic distance between the RIR-registered location and the observed location, enabling automated flagging of high-discrepancy addresses for further investigation.
- Match type classification — Records are annotated with a match type (e.g., rir_country_match) that characterizes the relationship between the registration metadata and the observed spatial data, enabling downstream consumers to apply appropriate confidence weighting.

Implications and Use Cases

Fraud Detection and Identity Verification

Applications that rely on IP geolocation to flag high-risk jurisdictions, enforce geographic access controls, or verify user identity are directly undermined by the tunneling scenarios documented here. A user in Shaanxi Province, China, or Banten, Indonesia, can present a Colombia or Chile IP address to any system that relies solely on IP geolocation — and that system will classify the connection as originating from a low-risk jurisdiction. IP Spatial provides an additional layer of verification that is not susceptible to this form of evasion.

Regulatory and Sanctions Compliance

Financial institutions, cloud providers, and other regulated entities use IP geolocation to screen for connections originating from OFAC-designated countries or other sanctioned jurisdictions. If an adversary routes traffic through a VPN endpoint in a compliant jurisdiction, standard IP geolocation will fail to detect the actual country of origin. IP Spatial-derived data provides the ground truth layer necessary to close this compliance gap.

Cybersecurity Threat Attribution

Accurate attribution of threat actor infrastructure requires knowing where network assets are physically located, not just where they are administratively registered. APT infrastructure frequently exploits IP address tunneling to obscure the country of origin of command-and-control (C2) traffic. IP Spatial data provides a spatial ground truth layer that can materially improve the fidelity of threat attribution analysis.

Content Licensing and Digital Rights Management

Media distribution platforms enforce geographic content licensing through IP geolocation. The well-known use of commercial VPN services to bypass geographic restrictions is an instance of the same structural failure demonstrated in this paper. IP Spatial can identify discrepancies between declared and observed location at scale, enabling more robust enforcement of territorial content licensing agreements.

Conclusion

The two case studies presented in this paper demonstrate a fundamental and systemic limitation of IP geolocation as it is currently practiced. Eleven commercially deployed geolocation providers, each drawing on RIR registration data, BGP routing tables, and associated heuristics, unanimously reported incorrect country-level locations for both IP addresses examined. In each case, the geodesic error exceeded 10,000 kilometers.

This is not a data quality problem that additional investment in traditional geolocation methodology can resolve. It is an architectural problem: IP geolocation infers location from administrative metadata that is structurally decoupled from physical device location in the presence of tunneling. Any system that relies exclusively on IP geolocation for location assurance must be regarded as having a critical, exploitable blind spot.

IP Spatial addresses this limitation at the root by replacing inference with observation. As the scenarios documented here illustrate, spatial telemetry-based location data provides both the accuracy and the confidence classification necessary to support high-assurance applications in fraud prevention, compliance, threat intelligence, and digital rights enforcement.

Appendix: Glossary of Key Terms

RIR	Regional Internet Registry. One of five global organizations (ARIN, RIPE NCC, LACNIC, APNIC, AFRINIC) responsible for the allocation and registration of IP address blocks within their respective geographic regions.
ASN	Autonomous System Number. A unique identifier assigned to a collection of IP routing prefixes under the control of a single administrative entity, used in BGP routing.
BGP	Border Gateway Protocol. The inter-domain routing protocol of the internet, responsible for exchanging routing and reachability information between autonomous systems.
IP Geolocation	The practice of inferring the geographic location of a network-connected device based on its IP address, using RIR registration data, BGP metadata, and heuristic probing techniques.
IP Spatial	A geolocation methodology and platform that derives device location from direct spatial observation of IP addresses via GPS-correlated network telemetry, rather than from registration metadata.
VPN (Virtual Private Network)	A technology that creates an encrypted tunnel between a device and a remote server, causing outbound traffic to appear to originate from the IP address of the remote server.
Geodesic Distance	The shortest path between two points on the surface of an ellipsoid (the Earth), used here to quantify the spatial discrepancy between RIR-registered and IP Spatial-observed locations.
match_type: rir_country_match	An IP Spatial record annotation indicating that the RIR-registered country and the geolocation-reported country are consistent, while the observed physical country differs. This is a field created during the research.
Confidence Level: HIGH	An IP Spatial classification indicating that the observed spatial data is supported by a sufficient volume of independent readings across multiple observation windows to meet the platform's threshold for high-confidence attribution.