Unmasking the Caller: Leveraging SIP Metadata, IP Analysis, and Network-Layer Verification to Combat VoIP Robocall Fraud

---

A Technical Whitepaper on SIP Protocol Forensics, GeoIP Attribution, and IP-Based Caller Identity for End-User Protection

---

Publication Date: May 2026

Version: 1.0 — Public Release

Audience: Telecom Engineers • Cybersecurity Professionals • Policymakers • Smartphone Platform Developers

Classification: Unclassified — For Public Distribution

Executive Summary

Voice over Internet Protocol (VoIP) robocall fraud has become one of the most pervasive forms of consumer crime in the United States, with Americans receiving **52.8 billion robocalls in 2024 alone** — averaging 1,627 calls per second. Fraudsters exploit a fundamental weakness in telephone infrastructure: the Caller ID system, which relies on the easily-spoofed SIP `From` header to display an arbitrary

phone number to the recipient, regardless of the actual call origin.

What fraudsters cannot hide, however, is the network-layer truth: **every SIP-based call carries the originating IP address in its Via and Contact headers**, as mandated by RFC 3261. This IP address — unlike the displayed phone number — reflects the actual source of the call and is structurally resistant to falsification without control of BGP routing infrastructure. When combined with GeoIP geolocation, WHOIS/ASN ownership lookup, and IP reputation intelligence, this IP data enables attribution of fraudulent callers to their true network source with high confidence.

This whitepaper documents a disproportionate and urgent targeting of **Asian-American and diaspora communities** — particularly Chinese-, Vietnamese-, Korean-, and South Asian-language speakers — by organized robocall fraud operations. In November 2025, the FBI issued PSA I-111325-PSA warning specifically of schemes where criminals impersonate both U.S. health insurers and Chinese law enforcement (公安), threatening victims with extradition and demanding payment under 24-hour video surveillance. These campaigns exploit cultural deference to authority, language isolation, and unfamiliarity with U.S. legal norms.

The STIR/SHAKEN cryptographic call authentication framework, mandated for U.S. carriers since 2021, represents important progress but remains severely limited by incomplete international adoption, TDM network gaps, and the exploitability of Gateway-level (C) attestation. Effective defense requires a multi-layer approach.

This whitepaper presents: **(1)** a forensic framework for SIP header and IP metadata analysis; **(2)** a composite IP-based caller risk-scoring model with weighted factors spanning GeoIP, ASN reputation, STIR/SHAKEN attestation, and behavioral patterns; and **(3)** a proposed smartphone display standard that surfaces IP-layer caller identity — including origin country, network owner, and risk badge — to end users in their preferred language, providing an accessible protection layer for communities most at risk.

Action is required from telecom carriers, smartphone platform developers (Apple iOS, Google Android), regulators, and law enforcement. The technical foundation exists today. What is lacking is coordinated implementation.

Table of Contents

Section 1: The Robocall Epidemic and the Asian Phishing Crisis

1.1 Scale of the Problem

The robocall crisis in the United States has reached a scale that defies casual description. According to the YouMail Robocall Index, U.S. consumers received **52.8 billion robocalls in 2024** — a figure that, despite representing a modest 4.5% decrease from 2023's 55.1 billion, still averages out to approximately **140.6 million calls per day**, or **1,627 calls every second**. Over the past four years, total annual robocall volume has persistently remained between 50 and 55 billion calls, demonstrating that incremental regulatory and technical countermeasures have thus far failed to achieve meaningful scale reduction. Each American adult is effectively subjected to approximately 200 robocall attempts per year.

The Federal Trade Commission (FTC) received over **2 million Do Not Call complaints** during fiscal year 2024, and the National Do Not Call Registry maintained over 253 million active registrations — a number that has long since ceased to deter

robocallers, who routinely ignore registry compliance. More alarmingly, the FTC's 2024 Consumer Sentinel Network data reveals that total reported consumer fraud losses surged to **$12.5 billion in 2024**, a 25% increase over the prior year. Government impersonation scams alone accounted for $789 million in reported losses — a $171 million year-over-year increase. The FBI's Internet Crime Complaint Center (IC3) recorded an even more sobering figure: in 2024, IC3 received 859,532 complaints representing **$16.6 billion in losses**, a 33% increase from 2023. Since the year 2000, IC3 has received more than 9 million complaints totaling over $50.5 billion in aggregate reported losses.

> ⚠️ **Key Statistic**
>
> The FBI IC3 2024 Annual Report documents that cyber-enabled fraud — including phone-based impersonation scams — accounts for
>
> 83% of all losses reported to IC3
>
> , totaling $13.7 billion in a single year. Call center scams alone generated 53,369 complaints and $1.9 billion in losses in 2024.

The robocall problem is not exclusively American. India processes hundreds of millions of fraudulent calls annually. Taiwan, South Korea, and Japan face persistent Chinese-language impersonation campaigns. Chinese diaspora communities in Canada, Australia, the United Kingdom, and across Europe are targeted by the same organized crime operations that exploit U.S. telecom infrastructure. The scale of transnational robocall fraud reflects a fundamental asymmetry: for a threat actor operating SIP trunking infrastructure in a permissive regulatory environment, the marginal cost of placing a robocall approaches zero, while the potential payout per successful victim can reach tens of thousands of dollars.

1.2 The Asian Phishing Epidemic

Within the broader robocall crisis, Asian-American and diaspora communities face a disproportionate and particularly sophisticated form of targeting. The vulnerability of

these communities is not incidental — it is systematically engineered by fraud operations that have invested in multilingual scripts, culturally-specific authority impersonation, and psychological pressure tactics calibrated to exploit the unique socio-legal anxieties of recent immigrants and non-native English speakers.

On November 13, 2025, the Federal Bureau of Investigation issued Public Service Announcement I-111325-PSA, specifically warning about an evolving financial fraud scheme targeting Chinese-speaking individuals residing in the United States. According to the FBI, criminals have been impersonating both U.S. health insurance providers and Chinese law enforcement (公安, *gong an*) officers in a coordinated two-stage attack:

1. **Stage One — Insurance Fraud Pretext:** The victim receives a call in Mandarin or Cantonese from a spoofed telephone number that appears to belong to a legitimate U.S. health insurer's claims department. The caller presents fabricated invoices — sometimes via screen-sharing through video communication software — alleging the victim has filed fraudulent medical claims.
2. **Stage Two — Law Enforcement Escalation:** When the victim denies the claim, the call is transferred to a second operator posing as a Chinese law enforcement officer. This impersonator demands personal identifying information, threatens the victim with extradition or prosecution in China, and demands a substantial cash payment as "bail." In extreme cases, the victim is instructed to remain connected via video surveillance for up to 24 hours — a form of virtual hostage-taking designed to prevent the victim from consulting family members or authorities.

> *"Scammers exploit*
>
> cultural nuances and mimic official procedures
>
> *to intensify the illusion of legitimacy." — FBI PSA I-111325-PSA, November 13,*

> *2025*

The FBI's explicit acknowledgment that these scams exploit "cultural nuances" reflects the structural sophistication of these operations. The threat of Chinese law enforcement action carries particular psychological weight for recent immigrants, who may have direct experience with authoritarian legal systems, who may have unresolved immigration status concerns, or who may simply lack the contextual knowledge to recognize that a Chinese law enforcement agency has no jurisdiction — and no legitimate communication channel — over a person residing in the United States.

The taxonomy of scam archetypes targeting Asian communities is broad and continues to evolve:

Scam Type	Target Community	Impersonation Vector	Pressure Mechanism
公安 (Gong An) Law Enforcement Scam	Chinese-speaking (Mandarin/Cantonese)	Chinese Ministry of Public Security, Interpol	Extradition threat, money laundering accusation
Health Insurance / Medicare Impersonation	Chinese-speaking residents	U.S. health insurer, Social Security Administration	Fraudulent invoice, account suspension threat
Fake IRS / SSA Scam (Vietnamese)	Vietnamese-speaking (Tiếng Việt)	IRS, Social Security Administration	Arrest warrant, tax debt demand
Fake IRS / SSA Scam (Korean)	Korean-speaking (한국어)	IRS, immigration authorities (USCIS)	Visa revocation, deportation threat
Parcel Delivery Scam	Chinese-speaking, South Asian	FedEx, DHL, customs authority	Customs fee demand, package confiscation
Wangiri / Missed-Call Callback	Pan-Asian (country-specific numbers)	None (social engineering via callback)	Premium-rate callback charges
Bank Account	South Asian	Major U.S. bank,	Account freeze,

Scam Type	Target Community	Impersonation Vector	Pressure Mechanism
Suspension Scam	(Hindi, Tamil, Telugu)	RBI impersonation	immediate fund transfer demand
Grandparent / Family Emergency Scam	Chinese, Korean, Vietnamese elderly	Family member in distress, attorney	Urgency, secrecy instruction

The IC3's 2024 Annual Report records that call center and government impersonation scams combined for over $2.3 billion in reported losses. While IC3 does not break out losses by ethnicity, law enforcement agencies with Asian community liaison programs — including FBI field offices in Los Angeles, New York, San Francisco, and Houston — have documented that elderly immigrants are among the most severely impacted demographic. The FBI IC3 data confirms that individuals aged 60 and older, while filing fewer complaints than younger age groups, experience dramatically higher per-incident losses — a pattern consistent with the high-value targeting observed in 公安 scams, where individual losses routinely exceed $50,000 and in some cases reach into the hundreds of thousands of dollars.

⚠️ Community Impact Note

Multiple factors compound vulnerability in immigrant communities: (1) language barriers that limit access to English-language fraud awareness resources; (2) cultural deference to authority figures, particularly uniformed officials; (3) unfamiliarity with U.S. legal norms — specifically, that no U.S. or foreign law enforcement agency collects bail payments by phone; (4) fear of immigration consequences that can make victims reluctant to report fraud; and (5) social isolation that limits access to informed friends or family who might interrupt the scam.

1.3 Why Current Defenses Fail

Despite years of regulatory action, technical standards development, and carrier investment in anti-robocall technology, the fundamental fraud problem persists. Understanding why requires examining each layer of the current defense posture.

Caller ID Spoofing Remains Trivially Easy on SIP Infrastructure: Any entity with access to a SIP trunk — which can be obtained from dozens of resellers for as little as a few dollars per month — can set the SIP `From` header to any phone number, including the direct lines of the Social Security Administration, the Chinese Embassy, or a specific individual's neighbor. Historically, many SIP trunk providers performed no validation of the presented number against the authenticated customer account.

STIR/SHAKEN Adoption Remains Structurally Incomplete: While the TRACED Act mandated STIR/SHAKEN implementation by major U.S. carriers by June 2021, the framework has critical coverage gaps. International calls — which constitute a large fraction of Asian-targeted impersonation calls — traverse international gateways where STIR/SHAKEN signatures are typically absent or downgraded to Gateway-level (C) attestation. Small VoIP resellers and rural carriers with weak Know Your Customer (KYC) processes remain exploitable origination points for domestically-signed fraudulent calls.

Consumer Awareness Is Critically Low in Target Communities: The FTC and FCC publish robocall awareness materials primarily in English. Multilingual fraud awareness campaigns targeting Mandarin, Cantonese, Vietnamese, Korean, Tagalog, and South Asian language communities remain insufficient relative to the scale of targeting. Elderly immigrants, who represent the highest-value targets for fraud operations, are among the least likely to have encountered anti-scam messaging in their native language.

Language-Specific Scripts Bypass Generic Detection: Most carrier-side and third-party robocall detection systems are trained on English-language call patterns and commonly spoofed U.S. phone numbers. Calls conducted entirely in Mandarin, targeting niche cultural anxieties specific to Chinese immigrants, are inherently less likely to be flagged by these systems. This creates a detection gap precisely in the vector most exploited by Asian-targeting fraud operations.

Section 2: How SIP Protocol Exposes Originating IP

2.1 SIP Protocol Fundamentals

The Session Initiation Protocol (SIP), defined in IETF RFC 3261 (Rosenberg et al., June 2002), is the dominant signaling protocol for Voice over Internet Protocol (VoIP) communications. SIP is a text-based, application-layer protocol modeled on HTTP and SMTP, designed to initiate, modify, and terminate multimedia sessions — most commonly voice calls — across IP networks. Unlike circuit-switched telephony, which establishes a dedicated physical circuit between parties, SIP-based calls are established through a message-exchange process that negotiates session parameters and then directs media to flow via the Real-time Transport Protocol (RTP) on separate logical channels.

> 📝 **Definition: Session Initiation Protocol (SIP)**
>
> SIP (RFC 3261) is a request/response protocol used to establish, manage, and terminate VoIP sessions. A SIP call begins with an INVITE request, proceeds through provisional responses (100 Trying, 180 Ringing), concludes session setup with a 200 OK and ACK exchange, and is terminated with a BYE message. Each message carries header fields analogous to HTTP headers, carrying metadata about the call including the addresses of participants, routing information, and authentication data.

The core SIP methods relevant to call fraud analysis are:

- **INVITE** — Initiates a call session; this is the primary message containing caller identity and routing metadata
- **ACK** — Confirms receipt of a final response to an INVITE (three-way handshake completion)

- **BYE** — Terminates an established session
- **CANCEL** — Cancels a pending INVITE before it is answered
- **REGISTER** — Registers a User Agent's location with a SIP registrar
- **OPTIONS** — Queries a server's capabilities; commonly used by robocall infrastructure for pre-call capability probing

Every SIP message — whether a request or response — consists of a start line, one or more header fields, an empty line, and an optional message body. For INVITE requests, the body typically contains an SDP (Session Description Protocol) payload that negotiates media parameters including codec preferences and, critically, IP addresses for media transport.

2.2 The Via Header: The IP Breadcrumb Trail

The Via header field is among the most forensically significant elements of a SIP message for fraud attribution purposes. Defined in RFC 3261, Section 20.42, the Via header serves the dual purpose of enabling response routing and creating an auditable record of every SIP proxy that handled the request. Per the RFC, each SIP proxy that forwards a request **must** append its own Via header value to the top of the Via header list before forwarding. The originating User Agent (UA) — whether a softphone, an Asterisk PBX instance, or a robocall autodialer — places the first Via header, which contains the IP address (or hostname) at which the UA expects to receive responses.

> 💡 **Forensic Key Insight**
>
> The
>
> bottom-most
>
> Via header in a received SIP INVITE (i.e., the first Via header as originally sent) contains the IP address of the originating User Agent — the actual machine that placed the call. This IP address cannot be fabricated without controlling the BGP routing of the IP block, making it structurally more reliable than any presented caller

identity information.

The following is a realistic annotated example of a SIP INVITE message as it arrives at a terminating carrier's Session Border Controller (SBC), showing multiple Via headers appended by intermediate proxies. This represents a typical fraudulent call claiming to originate from a U.S. Social Security Administration number while actually originating from a hosted SIP platform in Shenzhen, China:

```
INVITE sip:+12065550199@pstn-gw.us-carrier.net SIP/2.0 Via: SIP/2.0/UDP
203.0.113.47:5060;branch=z9hG4bK-us-carrier-01
<-- Terminating carrier SBC (US)
Via: SIP/2.0/UDP 198.51.100.22:5060;branch=z9hG4bK-intl-gw-hk01
<-- International gateway (Hong Kong)
Via: SIP/2.0/UDP 45.77.134.209:5060;branch=z9hG4bK-cn-proxy-sz01
<-- China proxy (Shenzhen datacenter)
Via: SIP/2.0/UDP 116.62.88.143:5060;branch=z9hG4bK-originator
<-- ORIGINATING IP (Alibaba Cloud, CN)
Max-Forwards: 66 To:
<
sip:+12065550199@us-carrier.net
>
From: "Social Security Administration"
<
sip:+18005721213@ssa.gov
>
```

```
;tag=a9f7k2
<-- SPOOFED: SSA's actual number
Contact:
<
sip:robocaller@116.62.88.143:5060
>
<-- Confirms originating IP
Call-ID: 9b3c1f27-4e22-4a8b-aed1-8f3d2c7e9b11@116.62.88.143 CSeq: 1 INVITE
User-Agent: FreeSWITCH-mod_sofia/1.10.6-release-5
<-- Gateway software fingerprint
P-Asserted-Identity:
<
sip:+85298765432@intl-gw-hk01.provider.net
>
<-- Hong Kong gateway's asserted ID
Allow: INVITE, ACK, BYE, CANCEL, OPTIONS, PRACK, UPDATE Supported:
100rel, timer, replaces, norefersub Content-Type: application/sdp Content-Length:
218  v=0 o=FreeSWITCH 1684923412 1684923413 IN IP4 116.62.88.143
<-- Confirms originating IP in SDP
s=FreeSWITCH c=IN IP4 116.62.88.143
<-- Media connection: same IP
t=0 0 m=audio 16394 RTP/AVP 0 8 101 a=rtpmap:0 PCMU/8000 a=rtpmap:8
PCMA/8000 a=rtpmap:101 telephone-event/8000 a=fmtp:101 0-16
```

Reading this INVITE from bottom to top in the Via header stack reveals the complete call path: the call originated from IP **116.62.88.143** (Alibaba Cloud, China region),

transited a Shenzhen-based proxy at 45.77.134.209, crossed an international gateway in Hong Kong at 198.51.100.22, and arrived at the U.S. terminating carrier through its SBC at 203.0.113.47. The `From` header falsely claims the call originates from the Social Security Administration's actual phone number — a complete fabrication. The `Contact` header and SDP `c=` line both corroborate the originating IP, providing independent confirmation.

NAT Traversal Considerations: In environments where the originating UA is behind a Network Address Translation (NAT) device, the Via header may initially contain the UA's private (RFC 1918) IP address rather than its public IP. The SIP `rport` parameter and STUN (Session Traversal Utilities for NAT) mechanism are designed to address this by enabling the receiving proxy to substitute the observed public IP. In practice, the first proxy in the chain — which receives the INVITE directly from the originating UA — always observes the true public IP of the NAT gateway, and this information is captured in server-side SIP logs even when not reflected in the header stack itself. Robocall farms operating at scale virtually always use cloud-hosted infrastructure with public IP addresses, making NAT-induced obfuscation rare in practice for industrial-scale fraud operations.

2.3 The From Header vs. P-Asserted-Identity

The SIP `From` header and the `P-Asserted-Identity` (PAI) header serve fundamentally different purposes and carry very different levels of trustworthiness — a distinction that is central to understanding how Caller ID spoofing operates.

The `From` header is set by the *calling party* — the User Agent initiating the call — and is propagated without modification through the SIP signaling chain. It represents a **claimed** identity, with no inherent authentication. On a standard SIP trunk with no number validation, the From header can be set to any E.164 telephone number or SIP URI whatsoever. This is the value that, absent STIR/SHAKEN processing, becomes the Caller ID displayed on the recipient's phone. Spoofing the From header is trivially accomplished in Asterisk with a single configuration line, in FreeSWITCH with an ESL command, or via any of dozens of commercial SIP spoofing APIs.

📝 **Definition: P-Asserted-Identity (PAI)**

The P-Asserted-Identity header (RFC 3325) is set by the

originating SIP service provider

— not the calling party — after the calling party has been authenticated to the provider's network. It represents the provider's assertion of the caller's true identity. PAI is a trusted network element: it is meant to be set by the carrier's equipment and stripped or re-set at trust boundaries. However, PAI is not cryptographically signed in basic SIP implementations and can be manipulated at gateways, making it more trustworthy than the From header but still not definitive without STIR/SHAKEN.

In the example above, the PAI header shows a Hong Kong gateway number — revealing that the international gateway processed the call from a source with a Hong Kong-format number, rather than the SSA number displayed in the From header. This PAI disclosure already contradicts the spoofed From identity and provides a forensic lead. In many robocall scenarios, fraudsters attempt to manipulate PAI as well, but this becomes increasingly difficult as the call traverses multiple carrier trust domains, each of which may re-assert PAI according to its own authentication records.

The related `P-Preferred-Identity` header (RFC 3325) is set by the UA to express a preferred identity to the originating carrier, which may then re-express this as PAI after validating the preference against the authenticated account. The `Remote-Party-ID` (RPID) header is an older, non-standard precursor to PAI used by some legacy Cisco equipment; it carries similar information but lacks standardized trust semantics.

2.4 The Contact Header and SDP Body

Two additional SIP message elements provide corroborating IP attribution data that cross-validates the Via header analysis.

The **Contact header** (RFC 3261, Section 20.10) specifies the direct URI — including IP address and port — at which the User Agent wishes to be reached for subsequent

requests within the dialog, such as ACK, BYE, or re-INVITE messages. Unlike Via headers, which accumulate proxy addresses as the message traverses the network, the Contact header reflects the UA's own direct address. In the example above, the Contact header reveals the same originating IP (116.62.88.143) as the bottom Via entry, providing independent corroboration. Discrepancies between Contact and Via headers can themselves be forensically significant, potentially indicating proxy misconfiguration or deliberate obfuscation attempts.

The **SDP message body** (Session Description Protocol, RFC 4566) contains the media negotiation parameters for the call, including the critical `c=` (connection) field, which specifies the IP address to which RTP media packets should be sent. In a direct call, this address matches the originating UA's IP. In calls using media proxies or TURN relays, the SDP may show a different address — but even this reveals the media infrastructure being used, which can be independently attributed. The `o=` (origin) field also typically includes the originating IP, as shown in the example above where FreeSWITCH recorded its own IP in both the o= and c= lines.

2.5 SIP Trunk Gateway Fingerprinting

Beyond IP address attribution, SIP metadata enables a more sophisticated form of fraud infrastructure identification: gateway fingerprinting. Robocall operations running at industrial scale — generating millions of calls per campaign — are not operated manually. They run on automated platforms, typically built on open-source PBX software such as FreeSWITCH, Asterisk, or Kamailio, often with commercial dialing front-ends. These platforms leave characteristic signatures across multiple SIP header fields that enable attribution of calls to specific fraud platforms even when IP addresses change.

SIP Header / Field	Fraud Indicator Pattern	Implication
User-Agent	FreeSWITCH-mod_sofia/ 1.10.x, Asterisk PBX 18.x, or custom strings like "sipstack/2.0"	Identifies the specific PBX software and version; shared across a fraud operation's infrastructure
SDP Codec Order	Non-standard codec	Platform-specific codec

SIP Header / Field	Fraud Indicator Pattern	Implication
	ordering (e.g., G.729 preferred over PCMU) or unusual codec sets	configuration fingerprint
Allow Header	Missing standard methods (e.g., no UPDATE or REFER) or unexpected methods present	Reveals specific SIP stack capability configuration
Supported Header	Specific extension combinations (e.g., 100rel+timer but no replaces)	Platform-version-specific SIP extension support
Call-ID Domain	Call-ID suffix matching the originating IP (e.g., @116.62.88.143) or a fixed internal hostname	Reveals true originating infrastructure even if other headers are manipulated
Branch Parameter Format	Non-random or patterned branch values (deviating from RFC 3261's z9hG4bK + random requirement)	Specific dialer software generates predictable branch patterns
RTP Port Range	Narrow or fixed port ranges in SDP (e.g., always 16384–16394)	Default FreeSWITCH or Asterisk configuration revealing unmodified platform settings
INVITE Rate	>10 INVITEs/second from a single IP; sequential or patterned ANI rotation	Automated dialing behavior incompatible with human-initiated calling
Max-Forwards Value	Fixed value (e.g., always 66 or 69) rather than the standard 70	Platform-specific default that may identify specific dialer software versions

The combination of these fingerprints enables fraud analysts to cluster calls across different campaigns, linking calls that use different spoofed numbers and different originating IPs to the same underlying fraud platform. When an IP changes — for example, as a fraud operator rotates through cloud infrastructure to evade IP-based blocking — gateway fingerprinting provides continuity of attribution that pure IP analysis cannot.

Section 3: IP Attribution — GeoIP, WHOIS, and Network Ownership

3.1 GeoIP Analysis

GeoIP analysis is the process of mapping an IP address to a geographic location using databases that correlate IP address ranges with physical locations. The two most widely used commercial GeoIP databases are MaxMind GeoIP2 and IP2Location, supplemented by the IPinfo.io API for real-time lookups. These databases are constructed by combining several data sources: Regional Internet Registry (RIR) allocation records, BGP routing announcements, active network probing, user self-reporting, and proprietary intelligence collection.

> 📝 **Definition: GeoIP**
>
> GeoIP is a class of technology that maps IP address blocks to geographic regions based on network registration data and behavioral intelligence. Accuracy varies significantly by geographic scope: country-level attribution is typically 95–99% accurate; region/state-level attribution achieves roughly 80% accuracy; city-level attribution is approximately 50–75% accurate and degrades further in rural areas and developing markets. Despite city-level limitations, country-level accuracy is sufficient to identify the fundamental geographic mismatch that characterizes most international robocall fraud.

For robocall fraud attribution, the critical application of GeoIP data is **geographic mismatch detection**: identifying calls where the claimed phone number's area code implies a U.S. location, while the originating IP resolves to a foreign country. A call displaying Caller ID "+1-202-456-1111" (spoofing the White House switchboard)

originating from an IP geolocated to Fujian Province, China, represents an unambiguous geographic mismatch with fraud probability approaching certainty.

GeoIP confidence scoring — offered by MaxMind GeoIP2 Insights and similar enterprise-tier services — provides a numeric confidence value (0–100%) for country and city attribution, along with additional metadata including:

- **Is Anonymous Proxy / VPN:** Flags IPs associated with commercial VPN services, which legitimate callers rarely use
- **Is Tor Exit Node:** Tor exit node traffic is inherently high-risk for any inbound call context
- **Is Hosting Provider:** Flags IPs registered to cloud hosting or datacenter organizations rather than residential or business ISPs — a critical indicator given that legitimate phone calls should not originate from AWS EC2 or Alibaba Cloud instances
- **Is Satellite Provider:** Satellite IP ranges may have low accuracy for country-level attribution and should be handled with reduced confidence

The "Is Hosting Provider" flag is arguably the single most actionable GeoIP attribute for robocall fraud detection. Legitimate residential or business callers connect to the PSTN through their local carrier's infrastructure, which operates from addresses registered to licensed telecommunications providers — not from EC2 instances, DigitalOcean droplets, or Alibaba Cloud ECS nodes. A call from a hosting provider IP should immediately trigger elevated risk scoring, regardless of the geographic location.

3.2 WHOIS and ARIN/RIPE/APNIC Ownership Lookup

While GeoIP provides geographic attribution, WHOIS lookups deliver organizational ownership data — identifying the registered entity that controls the IP address block from which a fraudulent call originates. The global IP address allocation system is administered by five Regional Internet Registries (RIRs), each responsible for a distinct geographic region:

Registry	Region Covered	WHOIS Server	API Access
ARIN	North America (US, Canada, Caribbean)	whois.arin.net	api.arin.net/rest
RIPE NCC	Europe, Middle East, Central Asia	whois.ripe.net	rest.db.ripe.net
APNIC	Asia-Pacific	whois.apnic.net	rdap.apnic.net
LACNIC	Latin America and Caribbean	whois.lacnic.net	rdap.lacnic.net
AFRINIC	Africa	whois.afrinic.net	rdap.afrinic.net

An ARIN WHOIS lookup for an IP address returns the **NetHandle** (the registered IP block), the **Organization** (the legal entity to whom the block is allocated), the organization's address, and in many cases contact information for abuse reporting. For IP addresses sub-allocated by cloud providers to customers, the lookup may return the cloud provider (e.g., "Amazon.com, Inc., Ashburn, VA") rather than the end customer — a finding that is itself informative, as legitimate telecom traffic should not originate from generic cloud allocation.

Complementing WHOIS data, **Autonomous System Number (ASN) lookup** provides the routing entity responsible for the IP. An ASN query for an IP returns the AS number (e.g., AS37963) and its registered name (e.g., "Alibaba Cloud (Beijing) Technology Co., Ltd."). The RPKI (Resource Public Key Infrastructure) framework allows verification of whether a given IP prefix has a valid Route Origin Authorization (ROA), confirming that the AS announcing the prefix is authorized to do so — a tool for detecting BGP prefix hijacking scenarios where an attacker fraudulently announces ownership of an IP range.

Consider the forensic implications of a real-world scenario: a call arrives claiming to be from the Social Security Administration (a U.S. government agency). The originating IP — extracted from the bottom Via header — is 116.62.88.143. An ARIN/APNIC WHOIS lookup immediately reveals this IP belongs to Alibaba Cloud Computing (Beijing), within the APNIC region. The ASN is AS37963, registered to Alibaba Cloud. No legitimate U.S. government agency routes telephony through Alibaba Cloud infrastructure in China. The mismatch is definitive.

3.3 Autonomous System Number (ASN) Reputation

Beyond identifying the owner of an IP block, ASN reputation databases provide historical intelligence about the fraud and spam history of that autonomous system as a whole. Just as email systems maintain sender reputation databases for domains and IP addresses, the telephony fraud intelligence community has developed ASN-level reputation profiles that enable rapid risk assessment of new IP addresses within a known-bad autonomous system.

> 📝 **Definition: "Dirty ASN"**
>
> A "dirty ASN" is an autonomous system that has persistent, documented associations with spam, fraud, or abusive network behavior across a significant fraction of its allocated IP space. Routing telephony through a dirty ASN is a strong fraud signal, particularly when the ASN has no legitimate telecom registration. Key dirty ASN intelligence sources include Spamhaus ASN DROP list, Cisco Talos IP and Domain Reputation Center, AbuseIPDB, and Team Cymru's IP-to-ASN mapping service.

The leading ASN reputation intelligence sources for telephony fraud assessment are:

- **Spamhaus BGP Ranking and DROP Lists:** Maintain lists of ASNs and IP blocks with the worst global reputations for spam and cybercrime origination; the DROP (Don't Route Or Peer) list identifies netblocks that have been "hijacked or leased by professional spam or cyber-crime operations."
- **Cisco Talos Intelligence:** Provides an IP and domain reputation center with scores derived from Cisco's global threat intelligence network, covering email spam, web threats, and network abuse.
- **AbuseIPDB:** A community-driven database of IP addresses reported for abuse, available via API for real-time lookup; useful for identifying IPs with recent fraud reports from other operators.

- **Team Cymru IP-to-ASN Mapping:** A free BGP routing intelligence service that provides ASN, prefix, and country information for any IP address, using BGP data from a globally distributed route collector network. Particularly valuable for real-time, automated enrichment pipelines.
- **Shodan:** A search engine for internet-connected devices; Shodan's database identifies IPs running SIP services (port 5060/5061), which can reveal open SIP proxies — common in robocall farm infrastructure — and their software versions and configurations.

3.4 Reverse DNS and PTR Record Analysis

Reverse DNS (rDNS) lookup — querying the PTR record for an IP address — provides an additional organizational identifier that complements WHOIS data. When a PTR record is configured, it typically reveals the hosting provider and, sometimes, the customer subnet. The pattern of the PTR record hostname is itself diagnostic.

Consider the contrast between the following PTR record patterns for a call originating from a given IP:

PTR Record Pattern	Interpretation	Risk Level
`atl-gw-01.att.net`	AT&T gateway in Atlanta — legitimate telecom infrastructure	**LOW**
`ec2-54-201-x-x.us-west-2.compute.amazonaws.com`	Amazon AWS EC2 instance — no legitimate telecom should originate here	**HIGH**
`116.62.88.143.static.cnchost.com`	Chinese datacenter static IP — non-telecom infrastructure in China	**CRITICAL**
`No PTR record (NXDOMAIN)`	Missing reverse DNS — common in ephemeral cloud instances and fraud infrastructure	**HIGH**
`voip-gw-3.legittelco.com`	Named VoIP gateway of a licensed carrier	**LOW**
`139.x.x.x.tencent-cloud.net`	Tencent Cloud instance	**CRITICAL**

PTR Record Pattern	Interpretation	Risk Level
	— Chinese cloud provider, non-telecom context	

The PTR record pattern for major cloud providers is highly predictable and can be pattern-matched in real time without a DNS lookup, using the GeoIP "Is Hosting Provider" flag or a locally-maintained list of known cloud provider IP ranges (CIDR blocks), which all major cloud providers publish publicly for firewall configuration purposes.

3.5 IP Reputation and Threat Intelligence Integration

The full power of IP attribution is realized only when individual data sources are integrated into a unified, real-time enrichment pipeline that can process each incoming call's originating IP and return a composite reputation score within the latency constraints of SIP processing (typically <200ms for a non-blocking pre-answer check).

The recommended integration architecture combines the following data sources in a locally-cached, regularly-refreshed threat intelligence database, supplemented by real-time API calls for dynamic data:

3. **GeoIP Database (MaxMind GeoIP2 Enterprise):** Local database lookup; <1ms latency. Returns country, region, city, ISP, organization, Is-Hosting-Provider, Is-VPN, Is-Tor, confidence score.
4. **ASN/WHOIS Data (Team Cymru / RIPE RDAP):** Locally cached ASN-to-organization mapping, refreshed daily. Returns ASN number, name, country, prefix.
5. **ASN Reputation (Spamhaus BGP Ranking, Talos):** Cached daily. Returns composite abuse score for the ASN.
6. **IP Reputation (AbuseIPDB API):** Real-time API call with 30-day cache. Returns abuse confidence score (0–100) and report count.

7. **PTR Record (Local DNS Resolver):** Real-time DNS lookup with aggressive caching. Pattern-matched against cloud provider signatures.
8. **Shodan Host Data:** Cached weekly. Returns open SIP ports, banner data, known vulnerabilities for the IP.
9. **STIR/SHAKEN Verification (STI-VS):** Real-time signature verification. Returns attestation level (A/B/C), origid UUID, certificate validity.
10. **Historical Call Database:** Internal database of previous calls from this IP, including first-seen timestamp, call volume, fraud reports from terminating users.

Section 4: Caller ID Spoofing — Mechanics and Limitations

4.1 How Caller ID Spoofing Works on SIP

The fundamental mechanism of Caller ID spoofing in SIP-based telephony is architectural rather than vulnerabilistic — that is, the ability to set an arbitrary From header is a designed feature of the SIP protocol, not an implementation error. SIP was originally designed for use within trusted network environments and was not conceived with the adversarial, open-internet context that characterizes modern VoIP deployments.

In a traditional SIP trunking arrangement, a business or individual contracts with a SIP trunk provider and connects to the provider via an IP-based SIP trunk. The provider authenticates the customer using SIP DIGEST authentication (username/password) or by IP address restriction. Once authenticated, the customer's PBX sends SIP INVITE messages with any From header value it chooses. Historically — and still today with many permissive trunk providers — the From header value is accepted without validation against the customer's authorized number set (Dial Number Identification Service, DNIS).

This means that a fraud operator who obtains a SIP trunk from a permissive provider can place calls that appear to originate from the IRS, the Social Security Administration, a local sheriff's department, or any private individual's phone number — simply by setting the From header accordingly. The displayed Caller ID on the recipient's handset reflects this spoofed value, with no indicator of inauthenticity visible to the user under current display standards.

The distinction between **presentation number** (what is displayed to the recipient) and **asserted identity** (what the network knows about the caller) lies at the heart of the spoofing problem. STIR/SHAKEN was designed to bridge this gap through cryptographic attestation.

4.2 STIR/SHAKEN: The Cryptographic Authentication Framework

STIR (Secure Telephone Identity Revisited) and SHAKEN (Signature-based Handling of Asserted information using toKENs) represent the industry's primary cryptographic solution to Caller ID spoofing. STIR is defined in a family of IETF RFCs, with the core protocol specified in RFC 8224 (Authenticated Identity Management in SIP) and the token format in RFC 8225 (PASSporT: Personal Assertion Token). SHAKEN is a specific ATIS-defined profile (ATIS-1000074) for how U.S. carriers implement STIR, defined in RFC 8588.

> 📝 **Definition: PASSporT (Personal Assertion Token)**
>
> A PASSporT is a signed JSON Web Token (JWT) embedded in the SIP
>
> Identity
>
> header. It contains the calling number (orig), the called number (dest), a timestamp (iat), and the attestation level (attest: A, B, or C). The token is signed using the originating carrier's private key, with the corresponding public key certificate retrievable from the STI Certificate Repository. Terminating carriers can verify the signature using the certificate to confirm that a legitimate, authenticated carrier vouches for the call's identity.

The SHAKEN framework defines three attestation levels that represent the degree of confidence a carrier can assert about the calling party's identity:

Level	Name	Meaning	Real-World Implication
A	Full Attestation	The carrier has authenticated the caller AND has authorized the caller to use the claimed telephone number. The carrier can fully vouch for the call's origin.	Highest trust. Carrier confirmed the customer is authorized to present this number. Spoofing at this level requires compromising a legitimate carrier account — still possible but significantly harder.
B	Partial Attestation	The carrier authenticated the customer's origination but cannot verify that the customer is authorized to use the specific telephone number in the From header.	Medium trust. Carrier knows where the call came from but not whether the number is legitimate. Common for business PBX customers with large number ranges.
C	Gateway Attestation	The carrier received the call from a gateway or upstream carrier that does not support STIR/SHAKEN (e.g., from a TDM/SS7 network or an international gateway). The carrier cannot vouch for the call's origin.	Lowest trust. Effectively equivalent to no attestation. Provides no meaningful authentication guarantee. Most international calls arrive with C attestation or no attestation.

The **origid** claim in the PASSporT is a UUID that uniquely identifies the origination source at the signing carrier. This identifier is the key to call traceback: when a fraudulent call is identified, the origid can be presented to the signing carrier (whose identity is known from the certificate) to demand identification of the specific customer or gateway that originated the call. This is the mechanism through which FCC-mandated traceback investigations proceed.

The STI (Secure Telephone Identity) certificate infrastructure involves multiple roles: the **STI-CA** (Certificate Authority) issues certificates to authorized carriers; the **STI-CR** (Certificate Repository) stores public certificates accessible for verification; the **STI-AS** (Authentication Service) at the originating carrier signs calls; and the **STI-VS** (Verification Service) at the terminating carrier verifies signatures. In the U.S., TransNexus and iconectiv serve as STI-CA authorities.

4.3 STIR/SHAKEN Limitations

Despite its conceptual soundness, STIR/SHAKEN's practical fraud prevention effectiveness is substantially constrained by a series of structural, technical, and compliance limitations:

International Coverage Gap: STIR/SHAKEN is mandated only for U.S.-registered voice service providers. International calls — including the majority of Asian-targeting fraud calls that originate from infrastructure in China, Southeast Asia, or the Middle East — arrive at U.S. international gateways without STIR/SHAKEN signatures, or are downgraded to C (Gateway) attestation by the receiving U.S. carrier. This creates a massive blind spot precisely for the call category most associated with Asian community targeting.

Exploitability of Legitimate Accounts: Fraud operators who obtain legitimate SIP trunk accounts with compliant carriers — using fraudulent business registration or stolen identities during the KYC process — can originate calls signed at Full Attestation (Level A). The cryptographic signature confirms only that a registered carrier customer placed the call, not that the displayed number or purpose is legitimate. Level A attestation from a KYC-negligent reseller is functionally equivalent to no attestation from a fraud-prevention standpoint.

TDM/SS7 Gap: A significant volume of calls still traverse portions of the Public Switched Telephone Network (PSTN) that operate on Time Division Multiplexing (TDM) and Signaling System 7 (SS7) protocols. SS7 predates modern cryptographic authentication and lacks any equivalent of STIR/SHAKEN. When a STIR/SHAKEN-signed call transits a TDM segment, its signature is lost — and the receiving carrier can only apply Gateway (C) attestation to the onward SIP leg.

Small Carrier and Reseller Compliance: While major carriers (AT&T, Verizon, T-Mobile) have fully implemented STIR/SHAKEN, hundreds of smaller rural carriers and VoIP resellers received FCC extensions on implementation timelines. Many of these remain partially compliant or sign calls without robust subscriber authentication, creating origination points that carry nominally valid STIR/SHAKEN signatures while lacking meaningful identity verification.

Consumer Comprehension: Even where STIR/SHAKEN functions correctly and carriers surface attestation-level information in call labeling ("Verified Caller" or "Spam Likely"), consumer comprehension of these labels — particularly among elderly immigrants and non-native English speakers — is minimal. A "Spam Likely" label does not explain why a call should be distrusted, nor does it provide the positive evidence of a fabricated foreign origin that IP attribution would supply.

4.4 FCC Enforcement and the TRACED Act

The Telephone Robocall Abuse Criminal Enforcement and Deterrence (TRACED) Act, signed into law on December 30, 2019, represented the most significant legislative action against robocall fraud to date. The TRACED Act directed the FCC to require voice service providers to implement call authentication technology (STIR/SHAKEN), to develop call blocking rules, and to establish a working group for call traceback. It also increased civil forfeiture penalties for illegal robocallers and extended the statute of limitations for enforcement actions.

The FCC's 2025 rulemaking (FCC 25-15) introduced additional requirements, including mandatory use of the SIP 603+ response code to inform callers when a call has been blocked by a carrier, providing transparency into call blocking decisions and enabling legitimate callers to address false positives. The FCC has also maintained and expanded Do-Not-Originate (DNO) lists — databases of telephone numbers (such

as Social Security Administration lines and IRS numbers) that should never legitimately appear as outbound Caller IDs, and that carriers are directed to automatically block or flag when presented as a calling number.

Despite these measures, the transnational nature of the robocall fraud ecosystem significantly limits domestic regulatory enforcement. The FCC has no jurisdiction over carriers or platforms operating outside U.S. borders, and international mutual legal assistance treaties (MLATs) are slow, resource-intensive, and rarely result in successful prosecution of overseas robocall operations. The practical enforcement gap means that regulatory pressure alone — without the technical countermeasures proposed in this whitepaper — cannot achieve adequate consumer protection.

Section 5: SIP Header Forensics and Metadata Analysis

5.1 Building a SIP Forensic Profile

A comprehensive SIP forensic profile is constructed from the complete set of header fields present in a SIP INVITE message. The following table documents all headers of forensic significance and their specific utility in fraud analysis:

SIP Header	RFC	Fraud Analysis Utility	Spoofability
Via	3261	Originating IP address (bottom entry); full proxy hop chain; NAT detection via rport	Low — network layer
From	3261	Claimed caller identity; mismatch with PAI reveals spoofing	Trivially spoofable
To	3261	Intended recipient; sequential number patterns reveal auto-dialing	N/A (recipient-set)

SIP Header	RFC	Fraud Analysis Utility	Spoofability
Contact	3261	Direct UA IP:port; cross-validates Via originating IP	Low — must be reachable
P-Asserted-Identity	3325	Carrier-asserted caller ID; more trusted than From; mismatch reveals gateway manipulation	Medium — trust-domain dependent
Identity	8224	STIR/SHAKEN PASSporT JWT; attestation level; origid for traceback	Low — cryptographically signed
User-Agent	3261	Platform fingerprint; identifies specific PBX/dialer software version	Easy to fake; consistency reveals lazy operators
Call-ID	3261	Globally unique call identifier; enables cross-carrier leg correlation for traceback; domain suffix often reveals originating IP	Low — must be unique
Date	3261	Originating system's timestamp; timezone reveals operational location; clock skew can indicate VPN/proxy use	Easy to manipulate
Allow	3261	Supported SIP methods list; platform fingerprint	Easy to fake; consistency reveals lazy operators
Supported	3261	Supported SIP extensions; platform fingerprint	Easy to fake

SIP Header	RFC	Fraud Analysis Utility	Spoofability
Max-Forwards	3261	Reveals number of hops already traversed; non-standard values fingerprint specific platforms	Automatically modified by proxies

Open-source tools enable real-time extraction and analysis of these headers at carrier-grade scale. **Kamailio** — a high-performance open-source SIP server capable of processing thousands of messages per second — supports scripting logic in its native Kamailio Routing Language (KRL) that can extract, log, and forward header data to external analysis systems. **FreeSWITCH** provides the Event Socket Layer (ESL) for real-time call event streaming. **Homer SIP Capture** (HOMER) is a purpose-built SIP analysis platform that uses the HEP (Homer Encapsulation Protocol) to mirror SIP traffic from production systems to a dedicated analysis stack, storing messages in PostgreSQL or Elasticsearch with full header parsing and a web-based analysis dashboard.

5.2 Metadata Fingerprinting Patterns

Robocall campaigns — unlike legitimate calling patterns — exhibit characteristic statistical behaviors in SIP metadata that enable automated detection and clustering even when individual calls appear syntactically valid. Machine learning techniques, particularly unsupervised clustering and anomaly detection, are effective at identifying robocall campaigns from SIP metadata streams.

Key behavioral patterns that distinguish robocall campaigns from legitimate traffic include:

- **INVITE Rate Anomalies:** A single IP address or subnet generating more than 10–15 SIP INVITEs per second vastly exceeds any legitimate business calling pattern. Industrial robocall operations routinely generate hundreds of calls per second from a single IP cluster.
- **ANI Rotation:** Automatic Number Identification (ANI) — the technical equivalent of Caller ID — is systematically rotated across campaigns to evade

per-number blocking. Robocall operations often rotate through sequentially numbered ANIs (e.g., +1-202-555-0100, +1-202-555-0101, +1-202-555-0102) or through lists of harvested legitimate numbers. The statistical distribution of ANI values from a given IP, when compared to expected distributions for legitimate business calling, reveals rotation patterns with high confidence.

- **Sequential Dialing (DNIS Patterns):** Robodialers often target sequential blocks of phone numbers (e.g., all numbers in a given area code prefix), producing a called-number (DNIS) distribution that follows a sequential rather than random pattern.
- **Short Answer Supervision Duration:** Robocall campaigns that use pre-recorded messages show characteristic call duration distributions — either very short (sub-10 second) durations when the call is not answered or the recipient hangs up immediately, or longer durations clustered around the length of the recorded script.
- **Infrastructure Reuse:** Despite IP address changes, gateway fingerprint clustering (User-Agent, codec order, SDP structure) enables linking calls across different IP addresses to the same underlying fraud platform — a form of campaign continuity detection.

Machine learning approaches applied to these features — including k-means clustering on User-Agent and codec fingerprints, DBSCAN for spatial clustering of originating IP subnets, and time-series anomaly detection on INVITE rates — can identify new robocall campaigns within minutes of their initiation, enabling proactive blocking before a campaign reaches its full scale.

5.3 The SIP Traceback Mechanism

The Industry Traceback Group (ITG), operated under the auspices of USTelecom, is the FCC-designated consortium responsible for implementing the SIP call traceback mechanism mandated by the TRACED Act. When a terminating carrier or law enforcement agency identifies an illegal robocall, the ITG initiates a traceback request that propagates hop-by-hop back through the call chain, requiring each intermediate

carrier to identify the upstream source of the call within 24 hours of receiving the request.

The ITG traceback process operates as follows:

11. **Initiation:** A traceback request is submitted to ITG, containing the Call-ID, timestamp, calling number, called number, and any available SIP header data from the terminating carrier's logs.
12. **Intermediate Carrier Notification:** ITG identifies the penultimate carrier (the carrier that delivered the call to the terminating carrier) using routing records and delivers a traceback request.
13. **Response Requirement:** FCC rules require carriers to respond to ITG traceback requests within 24 hours. Failure to comply can result in blocking by downstream carriers and FCC enforcement action.
14. **Recursive Tracing:** Each carrier that responds with an upstream source triggers a new traceback request to that upstream source, continuing until the originating carrier is identified.
15. **Enforcement Referral:** When the originating carrier is identified, ITG can refer the case to the FCC for enforcement, or law enforcement can subpoena the carrier for customer identification.

Key challenges to effective traceback include non-cooperative foreign carriers — particularly those in jurisdictions without mutual legal assistance treaty obligations to the U.S. — and VoIP resellers that provide minimal logging or that accept customers with inadequate KYC verification. Some fraud operations intentionally route calls through multiple resellers in quick succession to increase the depth of the hop chain and exhaust the 24-hour response window at each layer.

5.4 Real-Time SIP Monitoring Architecture

A carrier-grade SIP monitoring architecture for fraud detection combines network-layer packet capture with application-layer SIP parsing, indexing, and real-time alerting. The following architecture is representative of best-practice deployment at a mid-to-large scale U.S. carrier or SIP trunk provider:

Component	Technology	Function
SIP Traffic Mirror	HEP (Homer Encapsulation Protocol) agents on SBCs	Mirrors all SIP messages to the analysis stack without impacting production call processing; adds metadata (timestamp, source IP, agent ID)
SIP Capture Server	HOMER / Captagent	Receives HEP-encapsulated SIP messages, parses headers, indexes in database, provides web UI for fraud analysts
Message Store	Elasticsearch or PostgreSQL	Full-text indexed storage of parsed SIP messages; enables rapid historical search by IP, Call-ID, ANI, DNIS, or User-Agent
IP Enrichment Engine	Custom microservice (Python/Go)	Performs real-time GeoIP, ASN, PTR, AbuseIPDB lookups; attaches reputation metadata to each call record
Anomaly Detection	Elasticsearch ML / custom ML pipeline	Time-series anomaly detection on INVITE rates per IP; clustering on fingerprint features; sequential ANI detection
Alert & Response	Kibana Alerting / PagerDuty / custom SOAR	Real-time alerts to fraud analyst team; automated blocking rule generation for SBC ACLs
Fraud Analyst Dashboard	Kibana / Grafana	Real-time visualizations of calls/second per IP, geographic origin heatmap, STIR/SHAKEN distribution, ANI rotation patterns

Key operational metrics for the fraud monitoring dashboard should include: calls per second per IP (with configurable alert thresholds); ANI reuse rate (how many times a given calling number is used per unit time); geographic mismatch rate (percentage of

calls where GeoIP country differs from claimed area code country); codec anomaly rate (percentage of calls with non-standard codec configurations); and STIR/SHAKEN attestation level distribution (breakdown of A/B/C/none attestation across the call volume).

Section 6: IP-Based Caller Identity — A New Paradigm

6.1 Beyond Phone Numbers: IP as the True Identity Layer

The telephone number system, as currently constituted, is a presentation layer — a human-readable identifier that is mapped onto calls through a combination of carrier assignment and SIP header values. Phone numbers can be legitimately ported between carriers, reassigned to new owners, temporarily leased to businesses, and — as demonstrated throughout this whitepaper — trivially spoofed by any party with access to a SIP trunk. The phone number, by itself, carries no cryptographically verifiable link to a specific human caller or business entity at the moment of a call.

The originating IP address, by contrast, represents the **network-layer identity** of the call — the actual location in IP address space from which the SIP INVITE originated. While IP addresses can be reassigned through BGP routing changes, proxied through intermediary systems, or obscured by VPN services, this process requires either significant infrastructure control (BGP announcement authority over an IP block) or the use of intermediary services that are themselves attributable through their own IP addresses and organizational registrations. For the overwhelming majority of robocall fraud operations, which run on cloud-hosted or co-located server infrastructure with static IP allocations, the originating IP is effectively immutable across a fraud campaign and is tied to a specific, attributable network entity through WHOIS and ASN registration.

This white paper proposes a fundamental reconceptualization of caller identity: **the IP-layer identity should be surfaced alongside — and given equal or greater prominence than — the presented phone number**, as the former represents

verifiable network-layer reality while the latter represents a claimed identity that may have no relationship to the actual caller.

6.2 Components of IP-Based Caller Identity

A complete IP-Based Caller Identity (IP-CI) record for an incoming call should encompass the following components, assembled in real time from the sources described in previous sections:

Component	Source	Example Value	User-Facing Meaning
Originating IP Address	SIP Via header (bottom entry)	116.62.88.143	The actual network address of the call source
GeoIP Country	MaxMind GeoIP2	China (CN) — 98% confidence	"This call originates from China"
GeoIP City / Region	MaxMind GeoIP2	Hangzhou, Zhejiang — 72% confidence	Geographic specificity for display
ASN Name / Organization	Team Cymru / RIPE RDAP	AS37963 — Alibaba Cloud Computing (Beijing)	"Call routes through Alibaba Cloud"
WHOIS Registrant	APNIC WHOIS	Alibaba Cloud Computing Ltd., Hangzhou, CN	Registered owner of the IP block
Is Hosting Provider	MaxMind GeoIP2 / cloud range matching	TRUE	"Originating from a datacenter, not a telephone network"
PTR Record	Reverse DNS	116.62.88.143.static.cnchost.com	Network hostname confirmation
STIR/SHAKEN Attestation	SIP Identity header / STI-VS	None (no PASSporT present)	"Caller identity not cryptographically verified"
IP Abuse Score	AbuseIPDB	87/100 — 312 abuse reports	Quantified history of fraud from this IP

Component	Source	Example Value	User-Facing Meaning
ASN Reputation Score	Spamhaus BGP Ranking / Talos	HIGH RISK — significant abuse history	Network-level fraud association
First Seen / Call Volume	Internal carrier database	First seen: 14 days ago; 48,000 calls to date	Establishes campaign context and scale
Composite Risk Score	Risk scoring engine (Section 6.3)	94 / 100 — CRITICAL	Single actionable risk assessment

6.3 Risk Scoring Model

The following composite risk scoring model assigns a total score from 0 to 100, aggregating weighted factors across GeoIP, network ownership, STIR/SHAKEN attestation, behavioral indicators, and historical reputation. The score maps to four risk categories that drive distinct handling recommendations.

Risk Factor	Max Points	Scoring Logic	Example
GeoIP Country Mismatch	20	0: GeoIP country matches claimed area code country. 10: GeoIP country is neutral third country. 20: GeoIP country is high-risk origin (CN, RU, NG, NG, VN) mismatched against US area code claim.	Claimed US number (+1-202) resolving to China IP: 20 pts
Is Hosting Provider / Datacenter	15	0: IP registered to licensed telecom carrier. 8: IP registered to generic business ISP. 15: IP registered to cloud hosting provider (AWS, GCP, Azure, Alibaba Cloud,	Alibaba Cloud IP: 15 pts

Risk Factor	Max Points	Scoring Logic	Example
		Tencent, DigitalOcean, etc.).	
ASN Reputation Score	15	0: ASN with clean reputation, registered to a licensed telecom. 5: ASN with mixed reputation or non-telecom ISP. 10: ASN with elevated abuse history. 15: ASN on Spamhaus DROP or Talos high-risk list, or ASN with >75% abuse confidence.	ASN with 87% AbuseIPDB confidence: 13 pts
STIR/SHAKEN Attestation	15	0: Level A attestation (full, verified signature). 5: Level B attestation (partial). 10: Level C attestation (gateway only). 15: No STIR/SHAKEN signature present.	No PASSporT: 15 pts
IP Abuse History	10	0: No abuse reports. 3: 1–10 reports. 6: 11–50 reports. 8: 51–200 reports. 10: >200 reports or IP on known fraud infrastructure list.	312 AbuseIPDB reports: 10 pts
Sequential / Automated Dialing Pattern	10	0: INVITE rate <0.1/sec from IP; random DNIS distribution. 5: INVITE rate 0.1–2/sec; slightly elevated. 8: INVITE rate 2–10/sec or sequential	50+ calls/second from IP subnet: 10 pts

Risk Factor	Max Points	Scoring Logic	Example
		ANI/DNIS detected. 10: INVITE rate >10/sec or clear robodialer pattern.	
From/PAI Identity Mismatch	8	0: From and PAI consistent, matching GeoIP. 4: PAI missing (From only). 8: PAI present but contradicts From (e.g., From shows US number, PAI shows foreign gateway number).	From claims SSA, PAI shows HK gateway: 8 pts
Claimed Authority Number (DNO Match)	5	0: Called number not on any authority or sensitive list. 3: Number resembles but doesn't match a DNO list entry. 5: From header exactly matches a Do-Not-Originate list entry (SSA, IRS, FBI, government agency lines).	Spoofed SSA number on DNO list: 5 pts
Time-of-Day Anomaly	2	0: Call received during business hours in claimed geographic origin. 1: Marginal hours. 2: Call received at 2–5 AM local time of claimed geographic origin (midnight calling is a fraud indicator).	2:30 AM local: 2 pts

Total Score	Risk Category	Recommended Action	Display Color
0 – 30	**LOW RISK**	Allow call. Display standard Caller ID with green verification badge.	Green (#28a745)
31 – 60	**MEDIUM RISK**	Allow call with caution banner. Display IP origin country and organization. Log for review.	Yellow (#ffc107)
61 – 85	**HIGH RISK**	Display prominent warning with IP attribution details. Offer user option to decline. Flag for traceback.	Orange (#fd7e14)
86 – 100	**CRITICAL — LIKELY FRAUD**	Block call or route to voicemail with warning. Alert user with full IP attribution. Submit automatic traceback request.	Red (#dc3545)

Worked Example — Typical Chinese Law Enforcement Impersonation Scam:

Consider a call received at 14:30 PST (22:30 UTC) claiming to be from "Beijing Public Security Bureau Emergency Line" (+86-10-12345678), arriving at a U.S. terminating carrier:

- GeoIP: Originating IP resolves to Shenzhen, China; claimed number is a Chinese government number — same country. However, call is targeting a U.S. recipient from Chinese infrastructure with no legitimate US telecom path: **+15 pts** (high-risk origin targeting US resident without US number claim)
- Is Hosting Provider: IP registered to Tencent Cloud: **+15 pts**
- ASN Reputation: Tencent Cloud AS has elevated abuse history: **+8 pts**
- STIR/SHAKEN: No PASSporT (international call, no attestation): **+15 pts**

- IP Abuse History: 47 prior AbuseIPDB reports on this IP: **+6 pts**
- Dialing Pattern: 8 calls/second from the same /24 subnet: **+8 pts**
- From/PAI: No PAI header; From shows unverified foreign number: **+4 pts**
- DNO Match: Number not on US DNO list (foreign number): **+0 pts**
- Time Anomaly: 14:30 PST — business hours: **+0 pts**
- **Total: 71 / 100 — HIGH RISK**

For a variant of this scam where the caller also spoofs a U.S. number (e.g., a local police department number), add the GeoIP mismatch score of 20 and a DNO match of 3, bringing the total to approximately **94 / 100 — CRITICAL**.

Section 7: Smartphone Display of Real Caller Identity

7.1 The Case for On-Device IP Identity Display

Today's smartphone incoming call screen is a product of 1960s telephone network design thinking: it displays a phone number, sometimes a name from the local contact list, and — where carrier-side labeling has been implemented — a generic "Spam Likely" or "Scam Call" label. This display paradigm is entirely inadequate for the modern threat environment in which a call's presented phone number may be a government agency's legitimate published line, while the actual call originates from a cloud datacenter in Shenzhen controlled by a criminal organization.

The information asymmetry is stark: the fraud operator knows exactly what they are doing and why. The recipient — particularly an elderly immigrant with limited English proficiency — knows only what is displayed on their screen. Closing this asymmetry is not merely a technical improvement; for communities disproportionately targeted by 公安 scams and similar operations, it is an urgent public safety necessity.

The proposed IP-Based Caller Identity display standard would augment the existing phone number display with a structured panel providing:

- **Verified Phone Number:** The STIR/SHAKEN-attested number (if attestation is present), clearly distinguished from the unverified From header value if they differ
- **IP Origin Flag and Location:** Country flag icon + city/country name of the originating IP (e.g., "🇨🇳 Shenzhen, China")
- **Network Organization Name:** The WHOIS-registered owner of the originating IP block (e.g., "Alibaba Cloud Computing")
- **Risk Score Badge:** Color-coded badge (green/yellow/orange/red) with numeric score and category label
- **Attestation Level Indicator:** Visual indicator of STIR/SHAKEN attestation (Verified / Partial / Gateway / None)
- **Plain-Language Warning:** A human-readable statement in the user's device language, such as: "This call appears to originate from Alibaba Cloud in China. The displayed phone number may be spoofed. Chinese law enforcement cannot contact you by phone from a foreign datacenter."

7.2 Technical Implementation Pathways

Three distinct technical pathways exist for delivering IP-based caller identity information to smartphone users, each with different latency characteristics, data privacy implications, and implementation complexity.

Pathway 1 — Carrier-Side Enrichment (Recommended for Maximum Coverage): The terminating carrier performs all IP attribution lookups (GeoIP, ASN, WHOIS, IP reputation) during call setup, before delivering the call to the subscriber's device. The risk score and key IP attribution fields are encoded into the call delivery mechanism — specifically into the SIP INVITE delivered to the user's VoIP client, or into the push notification payload delivered via the carrier's IP Multimedia Subsystem (IMS) infrastructure to the device. Apple's CallKit and Google's Android telephony stack both support extended call information in incoming call notifications. This pathway requires carrier cooperation but provides the lowest latency and highest coverage, as it operates at the network layer before any app involvement.

Pathway 2 — Device-Side Lookup: The smartphone itself performs IP attribution lookups based on the source IP address disclosed during SIP signaling or observable from RTP/RTCP packet headers. This requires a privileged API for the device to access the originating IP of an active call — an API that does not currently exist in iOS or Android telephony frameworks but could be added. The device would query a cloud-hosted IP reputation API in real time (adding <500ms to setup time) and render the results in the incoming call UI.

Pathway 3 — Third-Party App Layer: Existing robocall screening apps (Hiya, Nomorobo, YouMail, Google Phone app built-in screening) already integrate with iOS CallKit and Android's Call Screening API to receive incoming call metadata and provide call disposition decisions and labels. These APIs could be extended to include originating IP address data (supplied by the carrier via enriched SIP delivery) for use by third-party screening services. This pathway leverages the existing ecosystem and requires the least platform-level change, but depends on carrier cooperation to provide the originating IP in a standardized field.

The technical integration points with platform-specific frameworks are as follows:

Platform	Framework / API	Current Capability	Required Extension
iOS	CallKit (CXCallUpdate)	Accepts caller name and "hasVideo" flag for incoming call UI customization; supports third-party Call Directory extensions for name/spam labeling	Add fields: originatingIPAddress, geoIPCountry, organizationName, riskScore, riskLevel, attestationLevel to CXCallUpdate
Android	Telecom API / Call Screening Service	CallScreeningService allows apps to inspect incoming calls and respond with allow/reject/silence; InCallUI extensions for	Add originating IP and attribution metadata to CallDetails bundle; extend CallScreeningService with IP attribution fields

Platform	Framework / API	Current Capability	Required Extension
		labeling	
Android (Google Phone)	Built-in Call Screening / Spam Detection	Google Phone app screens calls and displays "Suspected spam caller" with category	Integrate IP attribution data from carrier-side enrichment via CNAM/PAI extended fields; display in screening result
Carrier VoIP Apps	SIP/IMS Client	Display From header, PAI if available	Parse extended SIP headers carrying IP attribution data; render in in-call UI

7.3 UX Design for Multilingual Caller ID

The user experience design of IP-based caller identity must be calibrated for the communities most at risk — elderly immigrants with limited English proficiency, high visual impact requirements, and specific cultural context needs. Generic English-language fraud warnings are insufficient for a population whose primary threat vector is conducted in Mandarin, Vietnamese, or Korean.

The following UX design principles are proposed for the IP-based caller identity display:

Language Localization: All warning text must be rendered in the device's configured system language. For a device configured in Traditional Chinese (zh-TW) or Simplified Chinese (zh-CN), the warning "This call appears to originate from Alibaba Cloud in China" should render as: "此通话似乎来自中国阿里云。显示的电话号码可能是伪造的。中国执法部门无法通过外国数据中心的电话联系您。" This places the critical cultural context — Chinese law enforcement cannot contact you this way — directly in the language in which the scam is conducted.

Visual Hierarchy: The risk badge (color-coded) should be the most visually prominent element, larger than the phone number display for High Risk and Critical

calls. Country flag iconography provides an immediately interpretable geographic indicator without requiring text reading.

Accessibility: High-contrast display modes must be supported for elderly users with visual impairments. Font size for the risk label should scale with the device's accessibility font size setting. An optional audio announcement of the risk level ("Call from China, High Risk") could be provided for visually impaired users.

Progressive Disclosure: For Medium Risk calls, a compact summary (risk badge + country) should be shown by default, with a "Details" tap expanding to full IP attribution. For High Risk and Critical calls, the full attribution panel should be shown by default, as users should have this information immediately without requiring interaction.

Cultural Context Messaging: Beyond the generic warning, the system should maintain a library of culturally-specific scam context messages keyed to the combination of (claimed caller type, origin country, device language). For example: Claimed type "Chinese law enforcement" + Origin country "China" + Device language "Chinese" triggers the specific message: "中国公安机关不会通过视频通话要求您付款或保密。这是一个已知的诈骗手法。" ("Chinese public security agencies do not demand payment or secrecy via video call. This is a known scam technique.")

7.4 Privacy Considerations

The implementation of IP-based caller identity must be designed with rigorous privacy protections to prevent the creation of new surveillance infrastructure while deploying anti-fraud capability.

Data Minimization and Retention: IP attribution data assembled for an incoming call should be processed ephemerally — used to generate the risk score and display at call time — and should not be retained on-device beyond the active call session without explicit user consent. Carriers performing enrichment should log IP attribution data at the call level for fraud intelligence purposes, consistent with existing CPNI (Customer Proprietary Network Information) regulations, but should not associate enrichment data with the subscriber's communication history in a personally identifiable manner.

Consent-Based Crowdsourcing: Users who opt into a crowd-sourced fraud intelligence program can contribute anonymized, aggregate data about calls they report as fraudulent (IP block, ASN, risk factor pattern — not the specific call content or the user's identity) to shared threat intelligence databases. This enables rapid identification of new fraud campaigns across the carrier ecosystem. Contribution must be genuinely opt-in with clear disclosure, and contributing data must be stripped of any subscriber-identifying information before transmission.

CPNI Compliance: The IP address of the originating party is metadata about a communication and is subject to CPNI regulations under 47 U.S.C. §222. Carriers must ensure that IP attribution data disclosed to third-party screening apps through CallKit/Android APIs is covered by appropriate consent frameworks and data use agreements. Display of IP attribution to the called party is consistent with CPNI rules — recipients have a right to know about the nature and source of communications they receive.

International Compliance: For devices used by subscribers in jurisdictions subject to GDPR (e.g., European diaspora communities), IP addresses are classified as personal data. In these contexts, IP attribution data processing must be covered by a legitimate processing basis — the fraud prevention legitimate interest basis under GDPR Article 6(1)(f) is appropriate for this use case, provided a proportionality assessment confirms that the processing is necessary for the fraud prevention purpose and that appropriate safeguards are in place.

Section 8: Implementation Roadmap and Recommendations

8.1 For Telecom Carriers and SIP Service Providers

Carriers occupy the most strategically important position in the anti-robocall ecosystem — they are the network-layer entities that handle SIP signaling and are the only parties with complete visibility into both the originating IP and the presented caller identity for every call on their network. The following actions are recommended:

- **Implement Full STIR/SHAKEN with Strict KYC:** Ensure that all SIP trunk customers — including wholesale resellers — undergo identity verification equivalent to opening a bank account. Customers must prove authorization to use the telephone numbers they present in SIP signaling. Implement automated number validation that rejects SIP INVITEs where the From header number falls outside the customer's authorized number range.
- **Deploy Real-Time SIP Forensics Monitoring:** Implement HEP-based SIP capture with Homer or equivalent, feeding an IP enrichment pipeline that scores each call's originating IP against GeoIP, ASN reputation, and abuse databases in real time. Establish automated blocking rules for calls exceeding a risk threshold of 80+ points that do not have valid STIR/SHAKEN Level A attestation.
- **Enrich Call Delivery with IP Attribution Data:** Extend the SIP INVITE delivered to subscriber devices (or the push notification payload for VoIP subscribers) to include the originating IP, GeoIP country, ASN name, and composite risk score in standardized custom SIP headers (e.g., X-IP-Origin-Country, X-Risk-Score, X-ASN-Name) for consumption by device-side frameworks and third-party screening apps.
- **Participate in ITG Traceback Working Group:** Maintain complete, 24-hour-accessible call detail records (CDRs) with originating IP, Call-ID, ANI, and DNIS for a minimum retention period consistent with FCC traceback requirements. Respond to ITG traceback requests within the mandated timeframe and share intelligence on fraud campaigns with industry partners.
- **Contribute to Shared Fraud Intelligence:** Share anonymized IP-level fraud intelligence with ATIS, the i3 Forum, GSMA's Fraud and Security Group (FASG), and other anti-fraud consortia. Establish bilateral intelligence-sharing agreements with international carriers — particularly in Canada, the UK, Australia, and European markets facing the same Chinese-targeting fraud operations.

- **Establish Multilingual Consumer Fraud Hotlines:** Provide Mandarin, Cantonese, Vietnamese, Korean, and South Asian language support on fraud reporting and consumer awareness lines. Partner with community organizations serving these populations to distribute targeted fraud awareness materials.

8.2 For Smartphone Platform Developers (Apple and Google)

Apple iOS and Google Android collectively control the smartphone operating systems used by virtually all U.S. consumers. Platform-level changes to CallKit and the Android Telecom API would enable the most universal deployment of IP-based caller identity display, reaching all users regardless of carrier or third-party app:

- **Extend CallKit (iOS) and CallDetails (Android) with IP Attribution Fields:** Add structured fields for originating IP country, ASN name, risk score, and attestation level to the call information objects passed to incoming call UI rendering. These fields would be populated by carrier-provided data in the SIP delivery layer and by approved Call Directory / Call Screening extensions.
- **Develop Standardized Risk Badge UI Components:** Provide official platform UI components for displaying the risk badge, country flag, and organization name in the incoming call screen, ensuring consistent visual treatment across all apps and carriers. Avoid fragmentation where each carrier and app implements different, potentially confusing visual languages.
- **Require IP Attribution API in Carrier Certification:** Make the delivery of IP attribution data (at minimum, originating IP country and Is-Hosting-Provider flag) a requirement for carriers seeking CallKit or Android Telecom API certification. This creates a market incentive for carriers to implement the enrichment pipeline.
- **Expand Built-In Call Screening with IP Attribution:** Google's built-in call screening and Apple's Silence Unknown Callers feature should incorporate IP attribution signals. A call from a datacenter IP in China claiming to be a U.S. government agency should be automatically silenced or flagged, even without user action, as a platform-level default.

- **Partner with Threat Intelligence Providers:** Establish commercial partnerships with MaxMind, Team Cymru, AbuseIPDB, and similar providers to supply the real-time IP reputation data feeds needed for device-side or platform-side IP attribution lookups at scale.

8.3 For Regulators and Law Enforcement

- **Mandate IP Logging and Retention at Originating SIP Gateways:** The FCC should require all registered voice service providers to log and retain the originating IP address for every call they originate or transit, for a minimum of 18 months. This creates the evidentiary foundation for traceback investigations and civil/criminal enforcement actions.
- **Establish International SIP Traceback Frameworks:** The U.S. State Department and DOJ should prioritize negotiation of bilateral cybercrime cooperation agreements with India, China, the Philippines, and other source countries for robocall fraud, specifically addressing SIP traceback request compliance, evidence preservation, and prosecution of commercial robocall operators.
- **Pursue KYC-Negligent SIP Trunk Resellers:** FCC enforcement should prioritize actions against domestic SIP trunk resellers that provide service to fraud operations with inadequate identity verification, treating such negligence as a regulatory violation subject to meaningful financial penalties. Civil forfeiture authority should be used aggressively to seize revenues derived from fraudulent call traffic.
- **Fund Multilingual Consumer Awareness Campaigns:** Congress and the FTC should allocate funding for culturally-specific, linguistically-appropriate fraud awareness campaigns targeting Asian-American communities — specifically addressing the cultural context of 公安 scams, the impossibility of legitimate law enforcement contact via commercial VoIP, and the right to hang up without consequences. Partner with Chinese-language media, Vietnamese-language radio stations, Korean community centers, and South Asian community organizations for campaign distribution.

- **Establish Dedicated Asian Community Fraud Reporting Channels:** The FBI IC3 should establish Mandarin, Cantonese, Vietnamese, and Korean language reporting interfaces, with culturally trained intake staff who can communicate the reporting process effectively and provide appropriate follow-up guidance to traumatized victims.

8.4 For End Users

- **Never Trust Caller ID Alone for Authority-Claiming Calls:** No legitimate U.S. or foreign law enforcement agency, government department, bank, or healthcare provider will demand immediate payment by wire transfer, cryptocurrency, or gift card over the phone — and none will threaten arrest, extradition, or visa revocation if you do not comply immediately. If a caller makes any such demand, hang up immediately regardless of what Caller ID displays.
- **Verify Independently Using Official Published Numbers:** If you receive a call claiming to be from an organization — a health insurer, a government agency, a bank — hang up and call the organization back using the phone number printed on your insurance card, found on the official government website, or listed on the back of your bank card. Never use a callback number provided by the original caller.
- **Install a Call Screening App with IP Attribution Capability:** Use apps such as Hiya, Nomorobo, or Google Phone's built-in call screening, which incorporate carrier-side fraud intelligence. As IP attribution features become available (per the implementation pathways in Section 7.2), update to versions that surface originating IP information.
- **Report Suspicious Calls:** File reports with the FTC at reportfraud.ftc.gov and with the FBI IC3 at ic3.gov. Include any phone numbers, callback numbers, website links, and financial account details involved. Your report contributes to the shared intelligence that enables traceback investigations and fraud operation disruption.

- **Educate Vulnerable Family Members:** Elderly immigrants who are frequent targets of 公安 scams and similar operations need explicit, culturally-specific education from trusted family members. Specifically: explain that Chinese law enforcement has no jurisdiction in the United States; that no U.S. agency collects bail by phone; and that receiving a call does not create any legal obligation to comply.

Section 9: Conclusion

The VoIP robocall fraud epidemic — with its annual toll of 52.8 billion unwanted calls, $16.6 billion in reported cybercrime losses, and its particular devastation of vulnerable immigrant communities — is not an intractable problem. It is, in large part, an information asymmetry problem. Fraudsters know what they are doing. Their victims, and the systems meant to protect those victims, do not yet have access to the full picture that the underlying network infrastructure makes available.

As this whitepaper has demonstrated, the SIP protocol — the foundation of modern VoIP telephony — inherently exposes the originating IP address of every call in its Via, Contact, and SDP headers. This IP address is not a vulnerability to be patched; it is a feature of the protocol's design, intended to enable response routing. When properly surfaced and enriched with GeoIP geolocation, WHOIS organizational attribution, ASN reputation intelligence, and STIR/SHAKEN attestation data, this originating IP becomes a powerful fraud attribution instrument capable of revealing that a call claiming to originate from a U.S. government agency actually comes from an Alibaba Cloud server in Hangzhou, China — with quantified confidence scores and an audit trail suitable for law enforcement traceback.

The Asian phishing epidemic — the systematic targeting of Chinese-, Vietnamese-, Korean-, and South Asian-language-speaking communities in the United States — represents an urgent subset of this broader problem that demands culturally-informed, linguistically-accessible solutions. The FBI's November 2025 PSA I-111325-PSA made clear that 公安 scams, health insurance impersonation, and law

enforcement extortion calls are actively evolving and expanding. The communities most affected face a compounding set of vulnerabilities: language barriers that limit access to English-language fraud awareness, cultural deference to authority that makes impersonation particularly effective, and immigration-related anxieties that silence victims and suppress reporting. Technical solutions alone are insufficient — but technical solutions that surface IP-layer truth in the victim's native language, explaining that a call claiming to be Chinese law enforcement is actually routing through a commercial datacenter, can break the psychological architecture that makes these scams effective.

The risk scoring framework proposed in Section 6, the IP-Based Caller Identity display standard proposed in Section 7, and the implementation recommendations of Section 8 collectively constitute an actionable, technically grounded roadmap for materially improving caller identity integrity. The components are largely available today: GeoIP databases, ASN reputation feeds, STIR/SHAKEN infrastructure, Apple CallKit, and Android's Call Screening API all exist. What is required is coordinated action — carriers committing to IP enrichment pipelines, platform developers extending their APIs, and regulators creating the mandates and international frameworks that give the technical architecture legal force.

The phone call remains one of the most intimate and trusted channels of human communication. Restoring integrity to that channel — especially for those communities most systematically exploited when it is corrupted — is not merely a technical engineering objective. It is a matter of public safety, consumer protection, and social justice. The network-layer truth is available. The imperative now is to use it.

References and Bibliography

Standards and Technical Specifications

Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler. (2002, June). *RFC 3261: SIP: Session Initiation Protocol.* IETF Network Working Group. https://www.rfc-editor.org/rfc/rfc3261

Barnes, M., Jennings, C., Rosenberg, J., and M. Petit. (2006, November). *RFC 3325: Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks.* IETF. https://www.rfc-editor.org/rfc/rfc3325

Jennings, C., Peterson, J., and M. Watson. (2019, January). *RFC 8224: Authenticated Identity Management in the Session Initiation Protocol (SIP).* IETF. https://www.rfc-editor.org/rfc/rfc8224

Wendt, C., and J. Barnes. (2019, May). *RFC 8588: Personal Assertion Token (PaSSporT) Extension for Signature-based Handling of Asserted information using toKENs (SHAKEN).* IETF. https://www.rfc-editor.org/rfc/rfc8588

Peterson, J., and C. Wendt. (2019, January). *RFC 8225: PASSporT: Personal Assertion Token.* IETF. https://www.rfc-editor.org/rfc/rfc8225

Handley, M., Jacobson, V., and C. Perkins. (2006, July). *RFC 4566: SDP: Session Description Protocol.* IETF. https://www.rfc-editor.org/rfc/rfc4566

Alliance for Telecommunications Industry Solutions (ATIS). (2019). *ATIS-1000074: Signature-based Handling of Asserted information using toKENs (SHAKEN).* ATIS Standards. https://www.atis.org

Regulatory and Legislative References

U.S. Congress. (2019, December 30). *Telephone Robocall Abuse Criminal Enforcement and Deterrence Act (TRACED Act), Pub. L. 116-105.* 116th Congress. https://www.congress.gov/bill/116th-congress/senate-bill/151

Federal Communications Commission. (2025). *FCC 25-15: Advanced Methods to Target and Eliminate Unlawful Robocalls; Call Authentication Trust Anchor.* FCC Rulemaking. https://www.fcc.gov

Federal Communications Commission. (2021). *Report and Order on Remand: Implementing the TRACED Act; Call Authentication Trust Anchor.* FCC 20-136. https://www.fcc.gov

Law Enforcement and Government Reports

Federal Bureau of Investigation, Internet Crime Complaint Center. (2025, November 13). *PSA I-111325-PSA: Criminals Impersonate U.S. Health Insurance Providers and Chinese Law Enforcement to Target Chinese Speakers Residing in the United States.* FBI IC3. https://www.ic3.gov

Federal Bureau of Investigation, Internet Crime Complaint Center. (2025). *2024 Internet Crime Report.* IC3 Annual Report. https://www.ic3.gov/Media/PDF/AnnualReport/2024_IC3Report.pdf

Federal Bureau of Investigation, Internet Crime Complaint Center. (2024). *2023 Internet Crime Report.* IC3 Annual Report. https://www.ic3.gov

Federal Trade Commission. (2025, March 10). *New FTC Data Show a Big Jump in Reported Losses to Fraud to $12.5 Billion in 2024.* FTC Press Release. https://www.ftc.gov/news-events/news/press-releases/2025/03/new-ftc-data-show-big-jump-reported-losses-fraud-125-billion-2024

Federal Trade Commission. (2024, November). *National Do Not Call Registry Data Book for Fiscal Year 2024.* FTC Report. https://www.ftc.gov/reports/do-not-call

Federal Trade Commission. (2025). *Consumer Sentinel Network Data Book 2024.* FTC. https://www.ftc.gov/reports/consumer-sentinel-network

Industry Data and Threat Intelligence

YouMail, Inc. (2025, January 13). *U.S. Consumers Received Nearly 4.4 Billion Robocalls in December, 52.8 Billion in All of 2024, According to YouMail Robocall Index.* YouMail Robocall Index. https://robocallindex.com

MaxMind, Inc. (2025). *GeoIP2 Enterprise Database Documentation.* MaxMind Developer Documentation. https://dev.maxmind.com/geoip/docs/databases/city-and-country

Team Cymru. (2025). *IP-to-ASN Mapping Service Documentation.* Team Cymru Research. https://team-cymru.com/community-services/ip-asn-mapping

Spamhaus Project. (2025). *BGP Ranking and DROP List Documentation.* Spamhaus Technology. https://www.spamhaus.org/bgp

Cisco Talos Intelligence Group. (2025). *IP and Domain Reputation Center.* Cisco Systems. https://talosintelligence.com

AbuseIPDB. (2025). *AbuseIPDB API Documentation v2.* Marathon Studios. https://www.abuseipdb.com/api

Shodan.io. (2025). *Shodan API Documentation: SIP Device Search.* Shodan. https://developer.shodan.io

RIR and Network Registry References

American Registry for Internet Numbers (ARIN). (2025). *ARIN WHOIS and RDAP Service Documentation.* https://www.arin.net/resources/registry/whois

Asia-Pacific Network Information Centre (APNIC). (2025). *APNIC WHOIS Database and RDAP Service.* https://www.apnic.net/manage-ip/using-whois

RIPE Network Coordination Centre. (2025). *RIPE Database Query and RDAP Documentation.* https://www.ripe.net/manage-ips-and-asns/db

Internet Assigned Numbers Authority (IANA). (2025). *Autonomous System (AS) Numbers Registry.* https://www.iana.org/assignments/as-numbers

Open-Source Tools and Platforms

QXIP BV. (2025). *HOMER SIP Capture Server: Open Source VoIP Analysis Platform.* https://sipcapture.org

SignalWire, Inc. (2025). *FreeSWITCH Documentation: SIP Module (mod_sofia).* https://developer.signalwire.com/freeswitch

Kamailio Project. (2025). *Kamailio SIP Server Documentation: SIP Forensics and Security Modules.* https://www.kamailio.org/w/documentation

USTelecom — The Broadband Association / Industry Traceback Group. (2025). *Robocall Traceback Process and ITG Procedures.* https://www.ustelecom.org/the-industry-traceback-group

---

Unmasking the Caller: Leveraging SIP Metadata, IP Analysis, and Network-Layer Verification to Combat VoIP Robocall Fraud | Technical Whitepaper | May 2026 | Version 1.0 — Public Release

This document is provided for informational and educational purposes. All IP addresses, ASN identifiers, and operational examples in code blocks are illustrative. References to specific cloud providers are for technical attribution illustration only.