Deterministic Subscriber Attribution Under Carrier NAT

Closing the lawful-intercept and abuse-response gap that address sharing creates

An E^NAT IP briefing for network security, lawful-intercept, and regulatory-affairs leadership. For public release and evaluation.

Executive summary

Carrier-grade NAT solved IPv4 exhaustion and, in the same move, broke subscriber attribution. When one public address is shared by hundreds of subscribers, answering the question a subpoena actually asks — *which subscriber was using this address at this instant* — becomes an exercise in inference over a shared identifier. That inference is imperfect by construction: it loses precision, it raises the false-positive rate, and in documented cases it has implicated the wrong person. Practitioners now state the conclusion plainly — an IP address alone is no longer a reliable identifier.

There is a deterministic alternative. Instead of sharing a public address and reconstructing "who was who" after the fact from port-block logs, each subscriber is issued a unique private anchor inside the operator's own realm — assigned, not inferred — and that anchor is preserved *through* address translation rather than reconstructed after it. The result is per-subscriber attribution that is deterministic by design, survives the NAT hop, extends cleanly across wholesale and MVNO boundaries, and can be retained as evidence-grade record rather than best-effort log.

This brief sets out the gap, the regulatory exposure it creates, where existing approaches stop, and what a deterministic-anchor approach adds — including, candidly, what it does not do.

1. Why attribution breaks under CGNAT

Carrier-grade NAT (also large-scale NAT / NAT444) lets a single public IPv4 address serve hundreds of subscribers by allocating each a range of ports. A typical public address is shared by roughly 30 to 300 subscribers at any moment.

Attribution then requires mapping a subscriber to the public address *and the specific port* at an exact timestamp, and retaining that mapping — commonly for six months to two years to satisfy data-retention obligations. Two problems follow:

- **It is inference, not identity.** Because the public address is shared, attribution is a reverse lookup across a shared identifier. As the port-to-subscriber mapping is mined backward, precision falls and the false-positive rate rises. This is not a tuning problem; it is inherent to sharing the identifier.

- **The cost of getting it wrong is borne by the subscriber and the operator.** Where the reverse lookup resolves to the wrong party, an innocent subscriber can be implicated — a documented failure mode, not a hypothetical — with the operator carrying the liability and reputational consequences.

The operational burden compounds the risk: shared-port logging at carrier scale produces very large volumes of retained data that must be searchable years later, and the answer it yields is still probabilistic.

2. The regulatory and risk exposure

For the regulatory-affairs and legal reader, the exposure is concrete:

- **Lawful-intercept and data-retention regimes assume per-subscriber identifiability.** Frameworks such as CALEA in the United States, and their equivalents elsewhere, require an operator to identify the party using an address at a given time and to do so promptly. Shared-address attribution meets the letter through logging while degrading the reliability of the answer.
- **False attribution is a liability surface.** Every reverse lookup that resolves to the wrong subscriber is a potential wrongful action taken on the operator's data. The weaker the attribution primitive, the larger that surface.
- **Wholesale and MVNO arrangements widen the seam.** When subscribers of different operators share infrastructure, attribution must survive the handoff between administrative domains — precisely where shared-address inference is weakest.

3. What existing approaches do, and where they stop

None of the common responses is wrong; each stops short of deterministic per-subscriber attribution.

- **Deterministic NAT / port-block allocation** reduces logging volume by making the mapping reversible. But subscribers still *share the public address* — attribution remains at the granularity of an address plus a port range plus a time window, which is exactly where false positives live. It compresses the log; it does not remove the shared identifier.
- **IPv6** gives every device a unique address and ultimately dissolves the problem — but it does not address the IPv4-CGNAT reality carriers will operate for years through the transition, and it does not by itself solve attribution across wholesale/MVNO fabrics with overlapping private space.
- **Enhanced best-effort logging** improves retention and search, but the underlying answer is still an inference over a shared address.

4. A deterministic-anchor approach

The alternative changes the primitive rather than the logging around it.

- **Issue, don't infer.** Each subscriber or device is assigned a unique private anchor within the operator's own realm — unique even where subscribers' private address spaces collide. Attribution becomes a lookup of an identifier the operator *issued*, not a reconstruction across an address shared by hundreds.
- **Preserve identity through translation.** Coordinated translation at both ends of the boundary keeps the anchor invariant as traffic crosses it, so subscriber identity is carried *through* the NAT hop rather than reassembled afterward from shared state.
- **Answer "who, at time T" deterministically.** The anchor-to-public mapping is maintained as managed state, so the attribution query resolves to a single subscriber by construction — not by probabilistic port-block reversal.
- **Extend across owners.** Because uniqueness is maintained across multiple organizations on one fabric, attribution holds in wholesale, MVNO, and roaming arrangements, each association carried on the same issued anchor.
- **Optionally, make the record evidence-grade.** Above the attribution mechanism, the mapping can be retained as a signed, independently counter-signed record that verifies offline — turning "our logs indicate" into a record that neither the operator nor a counterparty could unilaterally rewrite. This evidentiary layer is an additional capability, distinct from the base attribution function.

The net effect is to move attribution from *inference over a shared identifier* to *lookup of an issued one* — removing the false-positive class at its root rather than compressing the logs that expose it.

5. What this does, and does not, do

Credibility depends on stating the boundary as plainly as the benefit.

It does: provide deterministic, per-subscriber attribution within the operator's realm; preserve that attribution through translation; extend it across wholesale/MVNO boundaries; and, at the higher tier, produce an evidence-grade, offline-verifiable record.

It does not:

- **Mandate itself.** Regulation requires attribution, not this mechanism; an operator can satisfy the letter of the law with port-block logging. This is a case for *better, lower-risk* attribution, not a compliance necessity.
- **Attribute conduct or content.** The mechanism establishes *which subscriber, over which carriage, at what time* — not what was done within an authorized session. It narrows the identity question; it does not read payloads.
- **Replace IPv6.** It is complementary to IPv6 migration, delivering deterministic attribution in the IPv4-CGNAT environment carriers are operating today and through the transition.

6. Why now

Three pressures are converging in 2026. Practitioner consensus now treats IP-based attribution as unreliable, raising scrutiny of exactly the reverse-lookup methods carriers depend on. Lawful-intercept and data-retention expectations continue to tighten around prompt, reliable per-subscriber identification. And automated, high-volume abuse increases both the number of attribution requests and the cost of getting any one of them wrong. The attribution primitive is becoming a risk decision, not just an engineering detail.

7. Provenance and availability

The approach is not conceptual. The mechanism — a coordinated dual-endpoint translation assigning a unique private anchor across overlapping address realms — is reduced to practice and embodied in the issued independent claims of two granted U.S. patents (US 9,680,792 and US 9,210,129; earliest priority 2 January 2013). It is available for evaluation as a managed service (MSL-as-a-Service) or as a licensed capability for integration into existing carrier NAT and lawful-intercept workflows.

Next step

A scoped evaluation against a defined attribution use case — for example, a wholesale or MVNO segment where shared-address attribution is weakest — will establish the deterministic-attribution and false-positive-reduction benefit on the operator's own traffic, with the evidence-grade retention tier assessed separately. We would welcome that conversation with your lawful-intercept, network, and regulatory-affairs teams.

E^NAT IP, LLC. This briefing is released for public evaluation. It describes a capability and its limits; it is not legal advice on any jurisdiction's lawful-intercept or data-retention obligations.