One IP is not one user.
Connection context tells you by how much.

An IP address describes network infrastructure. That is what keeps it working as user-level identifiers get harder to collect.

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

44%
of addresses show two or more distinct user agents

<0.2%
of domains resolve to a dedicated IPv4 address

Why addresses are shared

The internet ran out of IPv4 addresses years ago, and operators responded by pooling many subscribers behind shared public addresses. Research published this year by Stanford's Empirical Security Research Group, using measurement data from a global CDN, shows that sharing has increased every year for a decade.

Sharing takes two forms, and they distort measurement in opposite directions.

Shared at the same time

Carrier-grade NAT carries many unrelated devices behind one address at once. A frequency cap counted per address is exhausted by the first few, so the campaign stops serving before it reaches the rest.

Reassigned over time

On roughly 90% of shared addresses we observe few devices on any given day and many across a month. That is the signature of an address being reassigned rather than shared at that moment. A cap counted per address resets when the address moves, so the same subscriber can be served past the limit.

The assumption becomes a variable

Both forms distort CTV household graphs that join on the address and attribution models that credit a conversion to whoever sits behind it. Each of those systems is right to use the IP, and wrong whenever it treats one address as one identity.

Connection context turns it into a value you can read: how many devices sit behind an address, whether it is a household line or a carrier gateway, and whether it is also acting as a residential proxy or VPN exit node. Device count is the newest part of it, currently in private beta.

With a device count attached, a frequency-capping system treats a 200-device carrier gateway differently from a family broadband line, and a household graph weights its join by how shared the address actually is. None of this requires identity. Not a name, an email, a device graph or a consent chain. That is what makes network context useful where collecting or connecting user-level identifiers is increasingly constrained.

Sharing prevalence and user-agent figures: R. Habib, S. Singanamalla, M. Fayed and Z. Durumeric, On IP Addresses as Identifiers of Internet Users and Services, Stanford Empirical Security Research Group, 2026, using measurement data from a global CDN. Rotation share from IPinfo analysis of shared address behavior, February 2026. Device count is in private beta.

Test it against your own reach.

Run your frequency-capping or household-graph logic against a sample with device-count context attached, and see how much of your measured reach sits behind shared addresses.

Run A Sample →

ipinfo.io/adtech · sales@ipinfo.io