

September 2025

Geoff Huston

Geolocation and Starlink

"Where are you?" is not an easy question to answer on the Internet. The telephone system's address plan embedded a certain amount of physical location information in the fixed line network, and a full E.164 telephone number indicated your location in terms of your country, and your area within that country. The Internet did not adopt a geographic address plan which means that you are going to need a lot of additional information if you want to map an IP address into a location at the level of a country or a city.

Creating and maintaining such collections of geolocation data that maps IP addresses to a geolocation presents some challenges. Even a basic question, such as "How are you going to represent a location?" has a variety of answers. One could use latitude and longitude, but this has its own complications. What if you just wanted to map addresses into countries? You need a representation of a political map to translate these coordinates into a country. Or you could avoid these multiple layers of indirection and simply map IP addresses into countries. However, once you start referring to countries you run into a new set of questions, starting with the most basic ones of "What's a country?" and "What's a uniform way of naming them?" Thankfully these are not novel questions, and we can leverage the work of others to provide some answers here. There is the group that maintains the ISO 3166 standard, published by the International Organization for Standardisation that enumerates a list of codes of countries and dependant territories, using both 2 letter codes, three letter codes and three-digit numeric codes, all maintained by the imaginatively named "ISO 3166 Maintenance Agency," a group of 15 voting experts. There is also the United Nations Statistics Division, a body that maintains a number of related lists, including the "official" name of each country, a collection of numeric codes for each defined country, and a definition of a set of regions (groups of countries or groups of regions).

Why are these geolocation databases useful? There are obvious uses in the ongoing fight against various form of cyber-attacks, trying to de-anonymise the identity and location of the attacker. This information is also used in attempting to enforce various intellectual property rights that are often assigned to rights holders on a country-by-country basis. And then there are statistical reports. Countries like to compare themselves to others. But even simple questions, such as "How many Internet users are in each country?" are challenging to answer without the underlying seed data of a geolocation database.

There are a number of such IP address-to-location databases out there, but most are either private or only accessible on a subscription basis. In the research world many researchers have opted for the databases that are more generally available, and at APNIC labs for AP address-to-country mappings we rely on *Maxmind* and *ipinfo.io*. for this information.

The Regional Internet Registries also publish a two-letter country code in their IP number resource reports, and they are often used as a surrogate for IP-to-country mappings, but the data quality is low when assessed for quality as a source of geolocation information.

The reason why is based on differing assumptions between the data recording function of the RIRs and the needs of the geolocation function. The RIRs record the country of the principal office of the entity that was the recipient of the resource assignment. It does not record where these address and AS numbers are actually deployed on the Internet. In many cases, where an entity operates within a single country or economy, the RIR-recorded country code corresponds with the country where the associated addresses are being used. In other cases, where the IP addresses and AS numbers are used in other countries, the RIRs provide no indication that this is the case for these number resources. Also, the granularity of the data in the RIR registry is at a level of allocation, but when an assigned address block is divided up by an address holder and used in multiple countries there is no ability in the RIR data recording formation to track this internal subdivision and diverse deployment.

In general, it's not part of the RIRs' role to track where these number resources are being deployed. The RIRs' interest lies in accurately tracking "who" is assigned an address, and not "where" that address is being used.

Measuring ISP Market Share

At APNIC Labs we have been investigating if the collection of data that we have assembled as part of the measurement work can be used to track the ISP market share within each national economy. We are interested in trying to measure the effective level of inter-ISP competition within each national economy. The base of this derived competition measurement is a notional count of end users that are served by each ISP that operates in a national economy.

The measurement process starts with the estimated current population in each country. The data we use is sourced from the [United Nations Population Division](#). We use the mid-year population estimate from 2024 and apply the 2023-2024 growth rate to the period from mid 2024 to the present day to get an estimate of the current population of each country for this day.

The second data set we use is the proportion of the population of each country that are classed as Internet users. There are three possible sources for this data, the World Bank, the International Telecommunications Union (ITU) and the CIA World Factbook. We use the ITU data by preference, but the [three data sets](#) are well correlated in most cases.

The combination of these data sets gives us an estimate of the current Internet user population per country. It should be noted that this is not the number of “subscriptions” to a service, as it attempts to include the number of users behind each subscription. It also is supposed to avoid “double counting”, so where a user is part of a broadband service and also has a mobile service, then the user is still only counted once as an “Internet user”.

The third component of the data is the ad presentation data of the APNIC measurement program. We use Google Ads to deliver some [35M individual ad impressions per day](#). We use a geolocation database to map each user who received an ad impression to a country, and use a [local default-free BGP routing table](#) to also map each user to their "home" network. At this point we have now

assembled a set of "home" networks (origin AS numbers) and the geo-located country for each presented ad.

Assumptions

In this work we are making some pretty sweeping assumptions. These assumptions are somewhat questionable, but we've been forced to make them in the absence of generally available per-country data that is published by all countries in a timely and mutually consistent manner.

The first assumption is that Google's ad placement algorithms apply to all users within a given country uniformly. In defining the ad campaigns, we attempt to make the placement definitions as generic as possible, so that within each country the ad placements are roughly equivalent to a random sampling drawn from all users in that country. The implication of this assumption is that if an ISP has twice the number of users than another ISP in the same country, then its users will receive twice the number of ad impressions. This could be stated as: "The distribution of ad placement and the distribution of users across ISPs within any country are assumed to correlate."

The second assumption is that each user uses a single ISP for Internet access. This is not necessarily the case. For example, a user may use a local mobile service provider for their mobile Internet access and Starlink for their broadband access. We also have a user in their workplace using their workplace's ISP and using a consumer ISP when they are at home. These days many users have multiple mobile connections, and it is unclear how these multiple access methods correlate to ad placement, and through that to our measurements. The conclusion is that we can't account for such situations and in uniquely assigning each user to a single ISP in a country we tend to underestimate the user count for each ISP in consequence.

Due to the uncertainties that follow from these two major assumptions, the results we generate have an inevitable level of uncertainty. Some individual comparisons of this data against other sources where we have access to ISP market share data in individual countries point to an overall level of uncertainty of up to 15% or so in our estimates of users per ISP. Large consumer ISPs are still reported as having a large user population in the generated data, but the data for small networks is very uncertain.

The assumption of uniform distribution of ad placements across all ISPs within each country tends to fail where the number of placed ads in relation to the per-country user population is low. The best current example of this can be seen with the [Russian Federation](#), where ad placement in this country has plummeted since February 2023 (a consequence of the hostilities between the Russian Federation and the Ukraine and associated western sanctions being placed on Russia).

Another general assumption is that all users exist within a country. This does not hold for users on international flights using onboard Internet services, nor for ships at sea. In general, this factor should be insignificant for this exercise, given that as a proportion of the world's 4 billion users (or thereabouts!) this category of users is very small and should not distort the results to any significant extent beyond the already noted estimate of 15% uncertainty. But this general assessment does not hold when the ISP in question operates a service that is not constrained to any single country, such as a satellite-based service. Even so, when the satellite service operates as a wholesale service and provides connections as a service to ISPs, then this is not relevant to this form of measurement. If an ISP provides service in a country using IP addresses that are assigned to that ISP, then the conventional geolocation function will still provide usable results. The situation is different when the satellite operator provides its own retail services, using IP addresses that have been assigned to that satellite operator. This is the case for Starlink.

Starlink Data

The basic assumption here is that all IP addresses are used within a national realm. But this is not necessarily the case with users who are connected by a satellite service. What is the country when the IP service is provided to a ship on the high seas?

There are always exceptions to any generalisation, and some country views that are generated in this manner just stretch credibility too far.

Take Yemen, for example. A country with an estimated population of 10M people and 3.4M Internet users. The method described above gave the following result at the end of September:

Visible ASNs: Customer Populations (Est.)

Date: 22/09/2025

Rank	ASN	AS Name	CC	Users (est.)	% of country	% of Internet	Samples
1	AS14593	SPACEX-STARLINK	YE	6,233,929	59.22	0.147	321,186
2	AS30873	PTC-YEMENNET	YE	3,350,708	31.83	0.079	172,636
3	AS204317	ADENNET	YE	910,655	8.65	0.022	46,919
4	AS13335	CLOUDFLARENET	YE	28,647	0.27	0.001	1,476

Other Formats: [json](#) [csv](#)

Figure 1 – aspop report for Yemen from the 22nd September 2025

This measurement result is highly unlikely. It has been generated because over the past 60 days some 321,000 measurement advertisements originated from IP addresses that have been assigned to Starlink and Starlink's [geodatabase](#) geolocates these addresses to Yemen. The other three services providers appear to be the incumbent telco, Yemen Net, and a local ISP in Aden, Aden Net. The Cloudflare measurement is likely due to a combination of the local use of Apple's Private Data Relay and the Cloudflare's Warp product. Together, these three providers accounted for some 210,000 ad presentations over the same period. The result is that the algorithm we use assigned some 6M users in Yemen (or 60% of the country's Internet user population) to Starlink!

What factors might be at play here that would contribute to this anomalous result?

One potential factor is the volume of shipping in the Red Sea. These days it appears that the use of Starlink at sea is pretty much pervasive. A Starlink service is evidently a faster and cheaper communications service than that provided by Inmarsat and it has truly global reach. Given that the Starlink [geolocation](#) data attempts to map every Starlink IP address into one country or another, even ships at sea using Starlink get assigned an IP address that is mapped to some piece of land. Some 60 ships a day use the Suez Canal, and while the transit time from the Indian Ocean to the mediterranean sea is a few days, it's still a stretch to claim that shipping crew use of Starlink services alone accounts for some 50,000 ad impressions per day. These numbers imply that the use of Starlink by shipping may be part of the factors at play here, but it may not be the only contributory factor.

Another potential factor is that it's possible that Starlink's geolocation data does not reflect reality. The [Starlink availability map](#) indicates that Starlink has obtained national regulatory approval to operate in Yemen, Oman, Qatar, Bahrain, Israel, Jordan and Somalia, but not in Saudi Arabia, Egypt, Sudan, Eritrea, and Ethiopia. There have been persistent stories in a number of markets of Starlink resellers that set up a service in a country that has the necessary national regulatory approvals to use Starlink and then they ship the dish to a nearby location in a different country.

It's an open question as to the extent this is taking place, and if so then it's certainly plausible to guess that users in Saudi Arabia are using Starlink services that are registered in Yemen.

Does Yemen really have 6M Starlink users? That is extremely unlikely. How many Starlink users is the country likely to have? In neighbouring Oman, Starlink has a far more modest 0.08% market share, according to this same measurement technique. I would be surprised if the actual figure for in-country Yemen users is all that different. For the Yemen data, the high number might well be the result of a high count of Starlink-using passing maritime traffic being attributed to Yemen, and also some component of cross-country usage from perhaps Saudi Arabia and the United Arab Emirates, nearby countries where Starlink appears not to have local regulatory approval as yet.

Are there other countries with a similar problem of apparent over-representation of Starlink users? The ad placement data, assigned to countries using the Starlink [geolocation](#) data maps to 152 countries. In 21 instances, listed in Table 1, Starlink is used in more than 10% of the ad placement volumes, which looks to be somewhat questionable.

CC	Cover?	Ads	Est. Users	% Users	CC Name
SJ	Y	726	0	100%	Svalbard and Jan Mayen Islands
BL	Y	620	6,008	98%	Saint Barthelemy
TV	Y	7,980	5,799	92%	Tuvalu
KI	Y	42,234	1,7955	81%	Kiribati
PN	Y	16	19	72%	Pitcairn
YE	Y	321,673	6,256,291	59%	Yemen
NR	Y	6,864	4,071	56%	Nauru
CK	Y	16,220	4,802	50%	Cook Islands
MH	Y	7,857	7,805	34%	Marshall Islands
SS	Y	60,296	369,566	32%	South Sudan
MF	Y	1,412	4,468	24%	Saint Martin
VU	Y	214	22,423	22%	Vanuatu
NE	Y	140,318	1,076,585	21%	Niger
SD	N	348,986	3,517,776	19%	Sudan
TD	Y	78,690	292,985	17%	Chad
ZW	Y	311,093	801,754	15%	Zimbabwe
SB	Y	9,916	14,946	14%	Solomon Islands
MM	N	237,004	2,899,276	14%	Myanmar
FM	Y	9,824	6,164	14%	Micronesia
MG	Y	67,755	612,408	12%	Madagascar
TO	Y	4,881	5,304	11%	Tonga

Table 1 – Countries where Starlink attribution is greater than 10%

In the case of Svalbard other geolocation databases geolocate to Norway, whereas only the Starlink data set uses the SJ two-letter country code.

Saint Barthelmy, located in the Caribbean, is an overseas “collectivity” of France, with a population of some 9,000 people. Its former status was a commune as a part of Guadeloupe. While the Starlink geolocation database distinguishes between Guadeloupe and Saint Barthelmy, it appears that other databases do not draw a distinction between the two, hence the very high proportion of as placement in this country.

It is likely that the relatively high numbers of Starlink ad presentations in Tuvalu, Kiribati, Cook Islands. Marshall Islands, Saint Martin, Vanuatu, the Solomon Islands and Micronesia are due to shipping and yachting traffic. The relatively low GDP per capita in these island nations would tend

to indicate that Starlink services are unaffordable by such high percentages of the domestic population.

Starlink operates a Community Gateway service in Naru, and a traceroute to the IP address prefixes announced by this ISP (Cenpac, AS 5722) reveals a Starlink connection, presumably using inter-satellite laser link. The connections using Starlink's own IP addresses are presumably not part of Cenpac service, and these are likely to be an anomaly, presumably due to global roaming used by ships at sea. An examination of the routing table shows a similar community gateway has been deployed for the Tuvalu Telecommunications Corporation, Tamaani in Northern Quebec in Canada and the Federated States of Micronesia Telecommunications Corporation.

It's also possible that these additional ad placements could include an aircraft element, as there have been reports of Starlink selling a mobile access service to aircraft in flight, but as with ships at sea there is no published data on the uptake of this class of Starlink users.

There are a number of other anomalies in Table 1. Sudan and Myanmar both have a high ad placement rate, yet the [Starlink access map](#) indicates that the Starlink service is not available in either of these countries. If that is the case, then why does the Starlink geo data have IP address entries for both of these countries and why are so many ad placements being recorded from these IP addresses? In the case of Sudan, the Starlink gateway announcing these IP addresses is located in Mombasa in Kenya, and for Myanmar the relevant Starlink Gateway is located in Singapore. There are also high counts of ad placements for Starlink services that geolocate to Zimbabwe, Niger and Chad. The situation in the Cook Islands is potentially relevant here, where prior to regulatory approval to operate in the Cook Islands it was [reported](#) that domestic enterprises and some users were purchasing a Starlink service in New Zealand under a Roam Unlimited plan, and then shipping the equipment to the Cook Islands. There is no regulatory approval for Starlink to operate in South Africa, Namibia, Angola, and all of the countries in northern Africa and much of western Africa, and it's likely that there is a similar use of Starlink's roaming services to circumvent these local regulatory issues and purchase a roaming service elsewhere and use in in these countries.

For 20 of these 21 countries (the sole exception appears to be Pitcairn Islands) it's highly likely that the inferred level of use of Starlink within these countries is inflated by these factors, and the resultant view of the domestic ISP market is skewed as a result.

What is the role of Geolocation?

The rise of the use of satellite services for these global roaming services raises some basic questions about IP geolocation and its role.

Is this about the end user's precise physical location on the surface of the planet? Or is this about the national boundaries we've drawn on this surface, and assigning every user into one of these countries? In this case do we need to use a new geolocation code (or codes) for locations at sea? Is "at sea" defined by the conventional 12 nautical mile sea boundary? Or is some other interpretation of a margin where a country has a territorial sea claim?

What about ships in international waters? The conventional approach to ships at sea assert that the ship and its crew are subject to the laws of its flag state in international waters. What about aircraft in flight? It might appear that a similar situation to ships at sea may apply to aircraft in flight over international space, but a more commonly applied convention (the Tokyo Convention) is that the laws of the country of aircraft registration apply to an aircraft in flight for international flights irrespective of the location of the aircraft at any point.

So, what is the geolocation of the occupants of that ship or flight when accessing the Internet?

- Is it the location of the earth station used to pass packets to the spacecraft that relays the packet back to the ship or the aircraft?
- Or is it the location of the earth station used to receive the relayed packets and re-inject them into the terrestrial Internet?
- Should the geolocation be a reflection of the legal conventions of whose legal code applies to the ship or aircraft while travelling on an international journey.
- Should the geolocation reflect the country who is operating the satellite constellation that is performing the packet carriage function?
- In the case of ships travelling in territorial waters or aircraft flying over territorial airspace should the geolocation reflect the territory in question? In which case what should be do for those parts of the earth's surface that are not encompassed by accepted territorial claims?

There is a deeper assumption here concerning the behaviour of IP addresses. Does it even make sense to statically assign a geographic location to an IP address when the addressed device is in motion? What are the motivations for performing the location attribute assignment, and how can we implement the dynamic nature of such an assignment? There are no clear unambiguous answers to such questions, and perhaps that ambiguity reflects a common uncertainty that there is no clearly defined purpose for geolocation assignment in the first place.

APNIC Lab's Response to the uncertainties in Starlink's Geolocation

At APNIC Labs we've decided to override the Starlink geolocation data that refers to the 20 countries listed above and instead assign an “unclassified” designation to this part of the Starlink geolocation data.

It's not exactly a satisfying response to the problem, but it stops the distortion of the national measurements due to the increasing levels of usage of these satellite-based services for Internet access.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net