Anecdotal History of RFC 8805

Erik Kline
September 2025

A brief anecdotal history of how RFC 8805 came to be. Of no real technical value; purely for background. Opinions are my own.

Early IPv6 at Google

Google's first production IPv6 attempts used a dedicated domain, ipv6.google.com, to gain experience and minimize any adverse impact. This was first enabled during the "IPv4 outage" at IETF 71's Wednesday Plenary session (proceedings, ars technica) in March 2008.

Interest from Japanese network operators led to a second dedicated IPv6-supporting domain, ipv6.google.co.jp. While language and locale could often be inferred from domain names, this approach would not scale nor function correctly in all use cases.

Client IP addresses served as one language/locale hint, but enabling more public IPv6 services required an IPv6 geolocation database that could match IPv4 geolocation functionality. At this stage, IPv6 deployment was under 1% of Google users (statistics), and no commercial IPv6 geolocation data existed.

The fastest approach in these early days was to implement a configuration file mapping IPv6 prefixes to geolocation data. This was maintained by hand through professional connections with network operators (individual contacts via the IETF and various RIR/*NOGs).

IPv6 measurement experiment

During 2009/2010, Google began measuring global IPv6 adoption. This experiment ran on several Google websites and attempted to measure user connectivity to an IPv6 measurement site. In order to report broken or absent IPv6 connectivity, the experiment results were posted to an IPv4-only site (paper).

This enabled correlation between working IPv6 connectivity and IPv4 addresses used for reporting. Google implemented a production system to analyze these IPv6-to-IPv4 correlations and statistically estimate IPv6 prefix geolocations based on IPv4 geolocation data. Early challenges here included:

- statistically low volume of IPv6 samples and
- the complicating prevalence of IPv6 tunneling technologies (owing to the sparsity of native IPv6 deployment).

Edge deployments and ISP interaction

As Google's content cache edge deployments expanded, ISP coordination became crucial. This included devising a means for ISPs to inform Google of their geolocation mappings. This helped to address IP geolocation changes when a network operator relocated IP prefixes to a new in-country region but was also important for multi-national ISPs serving different nearby jurisdictions from adjacent prefixes.

Early geolocation data exchange focused on IPv4 prefixes, but for ISPs deploying IPv6 it was natural to extend the geolocation data exchange format to include IPv6 prefixes. An early version of the IP geolocation exchange format seemed like something worth bringing to the IETF, to raise awareness of the need for IPv6 geolocation data even though IPv6 deployment had only just reached 1%, by Google measurements (statistics).

Evolution at the IETF

In January of 2013 the first draft of what would become RFC 8805, "A Format for Self-Published IP Geolocation Feeds", was uploaded (history). At the time, it seemed like the geopriv working group might be a home for discussion. This working group, though, was nearing its completion. Some time was spent investigating the XML-based PIDF-LO work (RFC 5139). However, the Google geolocation engineering team preferred to continue working with their CSV-based format over an XML-based alternative.

Nevertheless, IETF review and feedback was helpful and the format was somewhat improved. The most notable change was the recognition that the "postal code" field SHOULD NOT be used for privacy reasons. This wasn't removed from the documented format, however; it was simply marked as deprecated with advice to not populate the field.

General privacy concerns remained, however, and the document did not have a clear path to progress within the IETF. In 2019, Warren Kumari noted the expired draft's utility for network operators, citing operational references and informal implementations. He suggested Independent Stream Editor publication and joined as coauthor to shepherd the document.

Post-8805 activities

Since publication, RFC 8805 "geofeeds" can be published and verified via RPKI mechanisms (RFC 9632), and are supported by some open source efforts (e.g. geofeed-finder, geolocatemuch). At least one study has been attempted to evaluate the deployment and utility of RFC 8805 "geofeeds" (ACM paper). Over the years, several individuals have reached out to the document authors about extensions, but any such work remains unfinished.

References

- IETF 71 proceedings; https://www.ietf.org/proceedings/71/

- Ars Technica article; https://arstechnica.com/features/2008/03/ietf-ipv6-switchoff/
- Google IPv6 Statistics; https://www.google.com/intl/en/ipv6/statistics.html
- Google IPv6 Measurement paper; https://research.google/pubs/evaluating-ipv6-adoption-in-the-internet/
- geofeed-finder; https://github.com/massimocandela/geofeed-finder
- geolocatemuch; https://geolocatemuch.com/
- Geofeeds ACM paper; https://dl.acm.org/doi/10.1145/3676869