The Internet Protocol Journal

September 2026 Volume 29, Number 2

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

You can download IPJ back issues and find subscription information at:
www.protocoljournal.org

ISSN 1944-1134

From The Editor

Deployment of *Internet Protocol Version 6* (IPv6) has been slow but steady since its initial standardization some 30 years ago. According to measurements performed by APNIC Labs, the world adoption rate now stands at somewhere around 41 percent. From time to time, new "simple" proposals collectively known as "IPv8" are put forward to address the initial motivation for IPv6, namely IPv4 address exhaustion. In our first article, Brian Carpenter explains why such proposals are a waste of time for everyone involved.

It used to be the case that a fixed-line telephone number could be associated with a geographic location right down to the street address of the telephone subscriber. With the advent of mobile telephones, *Number Portability*, and the introduction of *Voice over IP* (VoIP), it is obviously no longer the case. Last month, while attending a conference in Germany, I connected a vintage telephone to an *Analog Telephone Adapter* (ATA), which itself was connected via a mobile router to a *Session Initiation Protocol* (SIP) server in Sydney, Australia. It allowed me to use the vintage phone with its San Francisco-based telephone number in this new setting (see photo on page 41).

IP addresses are also used to determine geographic location of devices and users, but the accuracy of such location information depends on numerous factors. Last December, the *Internet Architecture Board* (IAB) organized a *Workshop on IP Address Geolocation*. Geoff Huston reports on the many approaches, problems, and solutions to IP geolocation that workshop attendees discussed.

All 97 previous issues of this journal are available on our website. You will find two ZIP archives; one will expand to a folder with each individual issue, and the other will expand into a single PDF file with more than 3,700 pages. Our website also contains index files for each volume as well as a cumulative index file of all issues to date. Check out **protocoljournal.org** for more details.

—Ole J. Jacobsen, Editor and Publisher
ole@protocoljournal.org

Why IPv6 Is So Complicated

by By Brian E. Carpenter, University of Auckland

There is no question that *Internet Protocol Version 6* (IPv6) is more complicated than IPv4, and people sometimes ask why. Surely it would have been much simpler to just add an extra 32 bits to the IPv4 address and change nothing else? In fact, every year or two people propose alternatives to IPv6 ("IPv8" is a generic name for such proposals, which involve mainly 8-byte addresses) because they have asked themselves that question. This article attempts to answer the question, and to show why such proposals are a waste of everybody's time, especially for the people who propose them.

The question has at least three possible answers:

1. Just adding bits to the address isn't as simple as it seems.
2. IPv4 was not the only network layer protocol in the world in 1994, and the others had good features that IPv4 lacked.
3. The IPv6 designers went "mad."

We will discuss each of these points in turn. To set the context, note that the decision to develop IPv6 rather than one of the other possible options was announced at the July 1994 *Internet Engineering Task Force* (IETF) meeting in Toronto, Canada. This announcement followed a long process that formally started at an *Internet Architecture Board* (IAB) workshop in 1991, as documented in RFC 1287[1] and led to an *Internet Engineering Steering Group* (IESG) report on future routing and addressing in late 1992, documented in RFC 1380[2]. Concerns about routing led to *Classless Addressing*[3,4] and *Border Gateway Protocol* Version 4 (BGP-4)[5] routing; concerns about address exhaustion led to agreement that we needed a new version of IP. But scaling up the routing and addressing systems was not the only concern in RFC 1380[2] (see point 2 above):

> **Although the catalog of problems above is pretty complete as far as the scaling problems of the Internet are concerned, there are other Internet layer issues that will need to be addressed over the coming years. These are issues regarding advanced functionality and service guarantees in the Internet layer.**
>
> **In any attempt to resolve the Internet scaling problems, it is important to consider how these other issues might affect the future evolution of Internet layer protocols.**

In any case, various proposals for the new version were drafted in 1991/93 and the IETF had no clear direction. A *Birds of a Feather* (BOF) session named "IPDECIDE" was held at the July 1993 IETF meeting in Amsterdam, The Netherlands. Its goal was to pick a direction, but the result was, um, indecisive.

Next, the IESG decided to set up an *IP Next Generation* (IPng) Directorate under two Area Directors—Scott Bradner and Allison Mankin—to support the decision process. This process led to the July 1994 choice, guided by RFC 1726[6], and explained in detail by RFC 1752[7] and the book *IPng: Internet Protocol Next Generation*, by Scott Bradner and Allison Mankin (editors), Addison-Wesley, 1995.[8]

Why Adding Bits Isn't Simple

IPv4 implementations, in 1994 and still today, have the 32-bit address format built into their code. Whether you expand the address size to 33, 64, or 128 bits, all IPv4 implementations will discard the packets. So, it's a matter of mathematical and physical fact that to expand the address size, you must change the protocol, and that means two things immediately:

1. You have to change the version number.

2. You have to add new code to handle the new version.

Furthermore, you don't want to split the Internet in two, so you must design a method of interworking between the old version and the new version. Annoyingly, you need to do that in a way that can be done completely in machines that know about the new version, because other machines don't know anything at all about the new version, by definition. So,

3. You need a coexistence technique so that updated systems, with the new protocol, can connect to old systems that know nothing of the new protocol.

Two minutes of thought show that this third requirement has only two solutions:

3A. Dual Stack, in which the new machines speak both the old (IPv4) and new (IPng) protocols.

3B. Translation, in which something translates addresses between the old and new protocols.

This reality has been known for more than 30 years, as outlined in RFC 1671[9], although people still sometimes try to deny it. To state it in graphical form, here is a diagram showing who can talk to whom, assuming the simplistic model of IPng with 64-bit addresses:

Versions	OLD	DUAL	NEW
OLD	32	32	XX
DUAL	32	64	64
NEW	XX	64	64

Protocol details, and the exact address length, don't matter. The XX cases can work only with protocol and address translation. All the complexities (and that really means *all*) of IPv4-IPng coexistence and transition are a result of this diagram, and the details of IPng design do not change this fact of nature.

Any purported design of a "better" or "simpler" IPng than IPv6 does not change this truth, however hard its authors try. In other words, the basic difficulties of IPv6 transition and coexistence have nothing to do with the design of IPv6.

Incidentally, "IPv8" proponents often ask why IPv6 didn't simply add some extra bits on the front of IPv4 addresses, instead of inventing a whole new format. Actually, we tried that: the "IPv4-Compatible IPv6 address" format was defined in RFC 3513[10], but it was deprecated by RFC 4291[11] because it turned out to be of no practical use for coexistence or transition. The related "IPv4-Mapped IPv6 address" format is still valid, and it has a role in the *Portable Operating System Interface* (POSIX) socket *Application Programming Interface* (API). Mappings of this kind also figured in the moderately successful coexistence technologies known as *6to4*[12, 13] and *Teredo*[14], which have now been overtaken by events.

Finally, it's worth remembering that IPv4 is itself no longer simple. Compared with IPv6 in 1994, we now have a whole lot of complications: subscriber-side *Network Address Translator* (NAT)[15,16,17], carrier-grade NAT, firewalls, *Internet Protocol Security* (IPsec)[18,19], *Virtual Private Networks* (VPNs), *Differentiated Services*[20], link-local addresses[21,22], *Content Distribution Networks* (CDNs), etc.

The Protocol Zoo

In the early 1990s, the Internet had not yet conquered the world (after all, the World Wide Web hardly existed before 1993), and there were many alternatives to IPv4 in use. Also, one of the most growly protocols in the zoo was wearing a smart business suit and lived in Switzerland with its rich friends—most governments and big businesses believed that the future network was certain to use the official international standard *Open Systems Interconnection* (OSI) protocol suite. To many people it seemed inevitable that an OSI network layer would replace IPv4. At the same time, numerous proprietary network layer protocols were in use. All the IETF had to offer were various competing IPng proposals, not even running code! All these other existing protocols had juicy features that IPv4 did not provide. Whatever IPng would be, it was expected to have at least some new features; plain IPv4 with bigger addresses was not what people expected or wanted.

Looking back, this situation was probably unfortunate, but it was a fact. IPng had to be better than IPv4, DECnet, Novell's *Netware*, etc.—and above all better than OSI.

Did the IPv6 Designers Go Mad?

That might be going a bit far, so the question should probably be: Did the IPv6 design suffer from *Second System Syndrome*, "...the tendency for a successful first system (often small and relatively elegant) to be followed by a second system that becomes over-engineered or bloated"?[23] Did we add stuff for its own sake?

The objective answer is "not much." First of all, IPv6 really is a conservative design—it doesn't change the basic IP model of connectionless packet switching with topological addresses. We added the flow label (harmless if unused, as it was for many years). We tweaked the fragmentation mechanism. We replaced IPv4 "options" with IPv6 "extension headers." (Just as IPv4 options are largely unused, so are IPv6 extension headers, except within limited domains.) Most important, we added *Stateless Address Autoconfiguration* (SLAAC)[24] and the closely linked *Router Advertisement* (RA)[25] mechanism, and the linked idea of an *Interface Identifier* (IID) as part of the address. SLAAC was inspired by *DECnet*, *Netware*, and *AppleTalk*, none of which required manual address configuration. SLAAC is partly redundant with *Dynamic Host Configuration Protocol for IPv6* (DHCPv6)[26], and the reason for that is that DHCP was new and unproven when SLAAC and RA were designed; DHCPv6 was actually a retrofit.

The author was too closely involved to say whether these changes and additions amounted to Second System Syndrome, but they were certainly not gratuitous changes. They caused very few of the problems during IPv6 deployment, and almost all of them come from the area of IPv4/IPv6 coexistence.

In addition, a rather political decision to require that IPv6 support IPsec, which was an immature technology at the time, caused some unnecessary confusion, which then slowed IPv6 deployment until the requirement was dropped after some years.

Do Not Make Things Worse

It's worth adding that some of the "IPv8" proposals over the years have included ideas that would make things worse rather than better—ideas like geographic addressing, address prefixes based on *Autonomous System Numbers* (ASNs), addresses with semantics encoded in their bits, and the list goes on. All these ideas, if implemented, would break Internet routing, make site renumbering even harder, risk running out of addresses yet again, simplify pervasive surveillance, or cause other forms of operational harm. (For one example, although geographic addressing has some obvious advantages when geolocation is wanted, it is incompatible with the way Internet interdomain routing works.)

It Would Take 25 Years Anyway

Getting IPv6 to about 50% deployment has taken more than 25 years. Any alternative or new proposal would be the same. Setting aside IPv6, consider that all of the following examples have taken decades, not years, to deploy Internet-wide:

1. Retiring *Frame Relay*
2. Replacing *Asynchronous Transfer Mode* (ATM)[44]
3. Deploying *Domain Name System Security Extensions* (DNSSEC)[45]
4. Deploying *Resource Public Key Infrastructure* (RPKI)[46]

Conclusion

The main reason for IPv6—and its only real reason for existence—had bigger addresses. The problems of coexistence were inevitable, and it was hard to find the best (or rather, least bad) solutions. Most of the difficulties of IPv6 implementation and deployment are not the result of the details of IPv6 design. Any address length greater than 32 would create all the coexistence and transition problems we have experienced since 1994. Both dual-stack deployment and translation (of protocol plus addresses) were mathematically inevitable. No alternative choice can possibly avoid these problems.

The community should think carefully before investing time and resources in such proposals.

Postscriptum

For the record, here are some of the proposals made over the years:

- Steve Deering, 1992, "The Simple Internet Protocol (SIP)," an early IPng candidate, had 8-byte addresses. SIP was assigned version number 6. It was only fully documented in 2018[27].
- Paul Francis, 1992, "The 'P' Internet Protocol (PIP)," an IPng candidate, was officially IPv8 for a while. It had variable-length addresses.[28, 29]
- Bob Hinden and Steve Deering, 1993/4, "Simple Internet Protocol Plus (SIPP)," an IPng candidate, had 8-byte addresses. It inherited the use of version number 6 from SIP. At the end of the IPng decision process, it mutated to 16-byte addresses and was the immediate precursor of IPv6.[30]
- **`draft-carpenter-aeiou`**[31] (1994)
- Jim Fleming touted "IPv8" and "IPv16" starting in 1996, but we have not found a coherent technical description of them. The best we have found is a cryptic statement on the *Africa Network Operators Group* (AfNOG) mailing list:

`IPv8 and IPv16 addresses are encoded in the **right-most 64-bits** of the 128-bit DNS. The left-most 64-bits are used for transition mechanisms.`

...and some archived diagrams:

https://web.archive.org/web/20000412004245/

http://www.unir.com/images/headers.gif

- **draft-terrell-ip-spec-ipv7-ipv8-addr-cls**[32] (1999)
- **draft-terrell-logic-analy-bin-ip-spec-ipv7-ipv8**[33] (2002)
- **draft-shyam-real-ip-framework**[34] (2014)
- **draft-omar-ipv10**[35] (2016)
- **draft-sambana-irtf-internet-protocol-sixteen**[36] (2022)
- **draft-thain-ipv8**[37] (2026)
- **draft-hause-asip**[38] (2026)
- **draft-subbiah-ipv7**[39] (2026)

Here's an interesting blog:
"What Is IPv8? Three Different Things With the Same Name,"
https://www.ip.network/blog/what-is-ipv8-protocol

References and Further Reading

[1] David D. Clark, Lyman A. Chapin, Vinton G. Cerf, Robert Braden, and Russel Hobby, "Towards the Future Internet Architecture," RFC 1287, December 1991.

[2] Phillip Gross and Philip Almquist, "IESG Deliberations on Routing and Addressing," RFC 1380, November 1992.

[3] Vince Fuller, Tony Li, Jessica Yu, and Kannan Varadhan, "Class less Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy," RFC 1519, September 1993.

[4] Yakov Rekhter, Robert G. Moskowitz, Daniel Karrenberg, Geert Jan de Groot, and Eliot Lear, "Address Allocation for Private Internets," RFC 1918, February 1996.

[5] Yakov Rekhter, Tony Li, and Susan Hares, editors, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, January 2006.

[6] Craig Partridge and Frank Kastenholz, "Technical Criteria for Choosing IP The Next Generation (IPng)," RFC 1726, December 1994.

[7] Scott Bradner, and Allison Mankin, "The Recommendation for the IP Next Generation Protocol," RFC 1752, January 1995.

[8] Scott O. Bradner and Allison Mankin, editors, *IPng: Internet Protocol Next Generation*, ISBN-13 978-0201633955, Addison-Wesley Professional, 1995.

[9] Brian E. Carpenter, "IPng White Paper on Transition and Other Considerations," RFC 1671, August 1994.

[10] Robert M. Hinden, and Stephen E. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture," RFC 3513, April 2003.

[11] Robert M. Hinden, and Stephen E. Deering, "IP Version 6 Addressing Architecture," RFC 4291, February 2006.

[12] Brian E. Carpenter and Keith Moore, "Connection of IPv6 Domains via IPv4 Clouds," RFC 3056, February 2001.

[13] Christian Huitema, "An Anycast Prefix for 6to4 Relay Routers," RFC 3068, June 2001.

[14] Christian Huitema, "Teredo: Tunneling IPv6 over UDP Through Network Address Translations (NATs)," RFC 4380, February 2006.

[15] Geoff Huston, "Anatomy: A Look Inside Network Address Translators," *The Internet Protocol Journal*, Volume 7, No. 3, September 2004.

[16] Kjeld Borch Egevang and Paul Francis, "The IP Network Address Translator (NAT)," RFC 1631, May 1994.

[17] Pyda Srisuresh and Kjeld Borch Egevang, "Traditional IP Network Address Translator (Traditional NAT)," RFC 3022, January 2001.

[18] Stephen Kent and Karen Seo, "Security Architecture for the Internet Protocol," RFC 4301, December 2005.

[19] Russ Housley, "Using Advanced Encryption Standard (AES) CCM Mode with IPsec Encapsulating Security Payload (ESP)," RFC 4309, December 2005.

[20] Steven Blake, Mark A. Carlson, Elwyn Davies, Zheng Wang, and Walter Weiss, "An Architecture for Differentiated Services," RFC 2475, December 1998.

[21] Stuart Cheshire, Bernard Aboba, and Erik Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses," RFC 3927, May 2005.

[22] Robert M. Hinden and Stephen E. Deering, "IP Version 6 Addressing Architecture," RFC 4291, February 2006.

[23] Wikipedia article on "Second-System Effect":
`https://en.wikipedia.org/wiki/Second-system_effect`

[24] Thomas Narten, Tatsuya Jinmei, and Susan Thomson, "IPv6 Stateless Address Autoconfiguration," RFC 4862, September 2007.

[25] Jaehoon Paul Jeong, Soohong Daniel Park, Luc Beloeil, and Syam Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration," RFC 8106, March 2017.

[26] Tomek Mrugalski, Bernie Volz, Michael C. Richardson, Sheng Jiang, and Timothy Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," RFC 9915, January 2026.

[27] Stephen E. Deering and Robert M. Hinden, editors, "Simple Internet Protocol (SIP) Specification," RFC 8507, December 2018.

[28] Paul Francis, "Pip Near-term Architecture," RFC 1621, May 1994.

[29] Paul Francis, "Pip Header Processing," RFC 1622, May 1994.

[30] Robert M. Hinden, "Simple Internet Protocol Plus White Paper," RFC 1710, October 1994.

[31] Brian E. Carpenter, "Address Extension by IP Option Usage (AEIOU),"Internet-Draft, Work in Progress, **draft-carpenter-aeiou-00**, 21 March 1994.

[32] Eugene Terrell, "Internet Protocol Specifications for IPv7 and IPv8 Address Classes," Internet-Draft, Work in Progress, **draft-terrell-ip-spec-ipv7-ipv8-addr-cls-06**, 8 December 1999.

[33] Eugene Terrell, "Logical Analysis of the Binary Representation and the IP Specifications for the IPv7 and IPv8 Addressing Systems," Internet-Draft, Work in Progress, **draft-terrell-logic-analy-bin-ip-spec-ipv7-ipv8-10**, 25 March 2002.

[34] Shyamaprasad Bandyopadhyay, "An Architectural Framework of the Internet for the Real IP World," Internet-Draft, Work in Progress, **draft-shyam-real-ip-framework-61**, 29 January 2020.

[35] Khaled Omar Ibrahim Omar, "Internet Protocol version 10 (IPv10) Specification," Internet-Draft, Work in Progress, **draft-omar-ipv10-13**, 21 October 2020.

[36] Bosubabu Sambana, "Internet Protocol version 16 (IPv16)," Internet-Draft, Work in Progress, **draft-sambana-irtf-internet-protocol-sixteen-01**, 9 February 2022.

[37] James Thain, "Internet Protocol Version 8 (IPv8)," Internet-Draft, Work in Progress, **draft-thain-ipv8-02**, 17 April 2026.

[38] Jordan Hause, "ASIP: AS-Structured Internet Protocol (128-bit)," Internet-Draft, Work in Progress, **draft-hause-asip-00**, 16 April 2026.

[39] Arunkumar Subbiah, "IPv7: Identity-Centric Network Protocol for Security, Proxy Mitigation, and Operability," Internet-Draft, Work in Progress, **draft-subbiah-ipv7-00**, 25 April 2026.

[40] Geoff Huston, "The Myth of IPv6," *The Internet Protocol Journal*, Volume 6, No. 2, June 2003.

[41] Geoff Huston, "IPv4 Address Depletion and Transition to IPv6," *The Internet Protocol Journal*, Volume 10, No. 3, September 2007.

[42] Geoff Huston, "The IPv6 Transition," *The Internet Protocol Journal*, Volume 28, No. 1, May 2025.

[43] Geoff Huston, "Internet Evolution," *The Internet Protocol Journal*, Volume 29, No. 1, May 2026.

[44] George Clapp and Mike Zeug, "Components of OSI: Asynchronous Transfer Mode (ATM) and ATM Adaptation Layers," *ConneXions—The Interoperability Report*, Volume 6, No. 4, April 1992.

[45] Miek Geben, "DNSSEC: The Protocol, Deployment, and a Bit of Development," *The Internet Protocol Journal*, Volume 7, No. 2, June 2004.

[46] Geoff Huston, "Resource Certification," *The Internet Protocol Journal*, Volume 12, No. 1, March 2009.

This article is adapted from the free open-source book about IPv6 called *book6*, written and maintained by a team of volunteers, who are all actively involved as users or providers of IPv6 services. You can find it, whether to read, comment, or contribute, at:

`https://github.com/becarpenter/book6/blob/main/Contents.md`

PDF and ePUB versions available—see "How to use this book."

BRIAN E. CARPENTER M.A., Ph.D., is an Honorary Professor in the School of Computer Science at Waipapa Taumata Rau, the University of Auckland, New Zealand. His research interests are in Internet protocols, especially the infrastructure layers including IPv6, as well as computing history. He was an IBM Distinguished Engineer working on Internet standards (1997–2007). Earlier he led the networking group at CERN, the European Laboratory for Particle Physics (1985–1996). He chaired the Internet Engineering Task Force (2005–2007), the Board of the Internet Society (2000–2002), and the Internet Architecture Board (1995–2000). E-mail: **`brian.e.carpenter@gmail.com`**

Our Privacy Policy

The *General Data Protection Regulation* (GDPR) is a regulation for data protection and privacy for all individual citizens of the *European Union* (EU) and the *European Economic Area* (EEA). Its implementation in May 2018 led many organizations worldwide to post or update privacy statements regarding how they handle information collected in the course of business. Such statements tend to be long and include carefully crafted legal language. We realize that we may need to provide similar language on our website and in the printed edition, but until such a statement has been developed here is an explanation of how we use any information you have supplied relating to your subscription:

- The mailing list for *The Internet Protocol Journal* (IPJ) is entirely "opt in." We never have and never will use mailing lists from other organizations for any purpose.
- You may unsubscribe at any time using our online subscription system or by contacting us via e-mail. We will honor any request to remove your name and contact information from our database.
- We will use your contact information only to communicate with you about your subscription; for example, to inform you that a new issue is available, that your subscription needs to be renewed, or that your printed copy has been returned to us as undeliverable by the postal authorities.
- We will never use your contact information for any other purpose or provide the subscription list to any third party other than for the purpose of distributing IPJ by post or by electronic means.
- If you make a donation in support of the journal, your name will be listed on our website and in print unless you tell us otherwise.

Geolocation in the Internet

by Geoff Huston, APNIC

Where on earth are you? Or, for that matter, where am I? I would like to look at the task of establishing the location of a device on the Internet, and how the Internet appears to know where I am.

To any owner of a mobile device, it seems odd to ask that question. With devices that have a *Global Navigation Satellite System* (GNSS)[1] receiver installed and enabled, where I am is an easy question to answer. If my device has an adequate view of the sky, it can calculate its position to an accuracy of within 5m or so these days. If I want to be more precise, I can use techniques such as *Real-Time Kinematic Positioning* (RTK)[2] to get to an accuracy of a centimetre or so by using a nearby reference point and specialised equipment. Even without centimetre accuracy, just knowing where I am, even within a meter or so, is great sometimes. It's great when I've attached a tag to my keys and then misplaced them! Or my luggage. Or when I want to hail a rideshare, or arrange a food delivery. Or when I'm lost!

However, location is also used in many other contexts. For example, on December 10, 2025, a new restriction[3] was applied in the country where I live, Australia, where people who are under 16 years of age cannot operate certain social media accounts. That raises two questions about me: What is my age? And am I in Australia? The first is a topic for a different time[4], but the second is another way of asking this location question. This time, the answer to that question does not necessarily demand a high level of precision, but the essential task for the social media operator is to map my location to a point on the geopolitical map of the world to see if I am within Australia's territorial boundaries or not. On my mobile device, with a clear view of the skies above, the *Global Positioning System* (GPS)[5] can do the job competently, but what about the location of my other devices, where GPS is not available? And in this case, it's my location that is important in order to comply with this legislation. The social media operator is required to know where I am, at least in the terms of "is my location within the territory of Australia or not?"

This activity of locating devices and users, termed *geolocation* or *geoloc*, was the topic of a workshop, hosted by the *Internet Architecture Board* (IAB) in early December 2025[6], and here I would like to relate my impressions from the discussions that took place in this workshop, as well as highlighting the material from subsequent discussions and presentations since then, notably at the first two *North American Network Operators' Group* (NANOG) meetings of 2026 on the same topic.

IP Geolocation

In the days of the fixed wire telephone network, an E.164[7] telephone number was assembled from a set of location-based fields. The full telephone number starts with a country prefix. All numbers starting with country code 61, for example, are located in Australia. The next set of digits is typically an *area code*. Although the size of this code varied between countries, regions within a country shared a common area code. Within an area there was often a discernible structure, so subscribers located in an area bounded by an individual *Telephone Exchange* (known as a *Central Office* in the U.S.)[8] would share a common intra-area code prefix. In addition, a line database was operated by the telephone company that had the precise location of every line termination point.

That database was not public data, but in an act that would be considered a breathtaking erosion of personal privacy today, the phone company annually published its *directory* of telephone subscribers, listing the subscriber's name, address, and telephone number.

In many ways the telephone system served as the model for the Internet in its guise as a public utility, and it raised the question: Could we use IP addresses in the same way we used telephone numbers, and derive a location from the IP address that a connected device is using? Is there a way to "map" an IP address to the location of the device that uses this IP address? The answer is not straightforward, as, unlike telephone numbers, there is no explicit hierarchy in IP addresses that would lend itself to derivation of even an approximate location. That does not mean that you can't use an IP address alone to derive a location, but you may need some form of a more detailed "map" that can translate an IP address to a location.

Early Geolocation Efforts

The motivation to adopt this standard format for mapping an IP address to a location was the increasing use of IPv6 in the public Internet some 20 years ago. At the time there were numerous sources of maps for geolocating IPv4 IP addresses, but no equivalent map existed at the time for IPv6 addresses. One of the early consumers of this location data was Google's search engine, which used the IP address of the querier to establish a location context for likely responses. For example, if a user queried for "pizza shop," Google's search system used the location information of the querier to trigger a search for "pizza shop near my location" and generated responses that matched both the content of the original query ("pizza shop") and its assumed context ("near my location").

Numerous IPv4 location maps were available at the time, typically distributed under commercial terms. These address maps were constructed from the base data published by the *Regional Internet Registries* (RIRs)[9], whose databases associated IP address prefixes with countries, and then further refined them with data gathered from other sources, including *ping* and *traceroute*, to infer a locational association between addresses.

No IPv6 location maps were executed at the time, and one way to rapidly bootstrap this IPv6 location data was to detect instances where a user had both IPv4 and IPv6 addresses, and the IPv6 address location could be associated with that of the IPv4 address.

However, this form of location mapping was little more than piecemeal guesswork. Informed guesswork to be sure, but it still lacked some form of ground truth. *Internet Service Providers* (ISPs) were encouraged to publish a map of the geolocations of the addresses they used in their network to allow others to regularly pick up that data and integrate the individual geofeeds into a larger geolocation map that covered the bulk of the IPv4 and IPv6 addresses in use at the time. The motivation for ISPs was that location information improved the experiences of their customers, with the ability of content providers to consume this location data to use default settings of locally used languages, to localise content, and to comply with relevant content regulations. Location information also aided in aspects of securing their users, such as detecting non-local attempts to access client accounts, for example.

However, the objective of feeding geolocation information into some system also had the associated problem of a common framework of describing a location. Latitude/longitude coordinates are one solution, but every client of such data also needs a set of maps that allow these coordinates to be translated into a civil address (country, stats, city, etc.). So we looked elsewhere for such a common specification of location.

Geofeeds

These days, the most common format for feeding data into a location "map" is specified in RFC 8805[10]. This RFC defines a simple text file, formatted as either *Comma Separated Value* (CSV) or *Extensible Markup Language* (XML) now, with each record describing a single address prefix with a set of location desccriptors: *ip_prefix, alpha2code, region code, city name,* and *postal_code*. The *alpha2code* is the ISO-3166[11] two-letter country code, and the *region code* is the ISO-3166-2 region code[12]. The code for the city name doesn't conform to any standard specification or language, because no such single standard exists, and the postal code is now deprecated. (In some countries, notably the United States, the full ZIP code is a 5-digit field for a city and a 4-digit extension that adds specific detail, such as a city block, or even a single entity if that entity receives a high volume of mail. The potential for high precision of postal codes creates a tension with personally identifying information, and for this reason postal codes have been deprecated from use in geolocation data feeds.)

It is now a common practice for network operators to publish this data about the location of their clients who have been assigned an IP address using this RFC 8805 format, and provide a pointer to the publication point of this data in their database record with the RIRs, using the *geofeed*: attribute of *inetnum* data objects in the RIR databases.

Another issue is that of knowledge and consent, particularly where specific location data is being used that may narrow down the set of clients encompassed by the location data. As RFC 8805 points out: "Operators who publish geolocation information are strongly encouraged to inform affected users/customers of this fact and of the potential privacy-related consequences and trade-offs." There is also a follow-up document, RFC 9632, "Finding and Using Geofeed Data,"[13] that specifies how to augment the *Routing Policy Specification Language* (RPSL)[14] *inetnum*: class to refer specifically to these RFC 8805 geofeed files by URL. This specification also describes an optional scheme that uses the *Resource Public Key Infrastructure* (RPKI)[15] to authenticate these geofeed files. The use of RPKI is interesting, as it is a response to address an aspect of authenticity of geolocation information. The concept behind these geofeed signatures is that the address holder can use their private key to sign a geolocation entry, asserting that they have some legitimacy in attesting to the location of that address, as well as allowing the user of the file to validate the currency completeness of the report as an exact replica of the originally published data. However, it does not directly address the veracity of the location information. In other words, I can still deliberately lie about the location of an IP address, but if I sign it with my key then I can't deny that it was my lie!

This work was informed by the *geopriv* Working Group[16] of the *Internet Engineering Task Force* (IETF), a group that focused on defining standards for handling geographic location information securely and privately within Internet protocols, creating frameworks for location-aware applications, and real-time communication while keeping the development of core location discovery technology out of scope. The working group published RFC 6280, "An Architecture for Location and Location Privacy in Internet Applications,"[17] and RFC 5870, "A Uniform Resource Identifier for Geographic Locations ('geo' URI),"[18] establishing common formats, APIs, and protocols, like *HTTP-Enabled Location Delivery* (HELD)[19] for location sharing, helping ensure privacy was integral to location data management.

Of course, location has many other implications beyond figuring out where nearby pizza joints are. Harking back to Google's search engine, some results are not permissible in certain countries, so searches for any sensitive terms would be affected. This situation does not necessarily prevent a multi-national service platform, such as Google, from assembling pointers to such content, but is does restrict the presentation of such pointers to end users based on the assumed location of the user. There is also the consideration that many countries have quite unique web cookie management policies, entailing a customisation of application behaviour to match local policies based on the assumed location of the user's device. Many people, including myself, however, find it oddly bizarre that we are asked for specific directions about management of these cookies. For example, when I visit certain web sites popups appear, and I am uncertain as to what precisely they are asking for, why, or what granting or denying such permission may imply for my experience as a service client or even my personal privacy!

It seems to me that these kinds of popups are a classic misapplication of the concept of *user consent*, where I am asked to provide consent for an action over which I have no competent understanding related to what such consent implies!

So far, I have described the exercise as one that relates to my location by assuming that it is the same as the location of the device I am using. I have various forms of remote device access tools that allow me and my device to be arbitrarily distant, and in general the assumption behind my location through my device location ignores the possibility that these two locations might be distinct.

Let's also consider "teleportation" tools, such as *Virtual Private Networks* (VPNs)[20,21] and similar forms of relay tools[22] and proxies, where the location of interface to the public Internet is different from the location of my device. To any onlooker or service provider, my device appears to be located at the tunnel egress point to the extent that its appearance is indistinguishable from any other device that is physically located at the same tunnel egress point.

What appears to be lacking here is a consistent ontology about the meaning of "location," relating to the location of network infrastructure equipment that acts as a tunnel endpoint, the location of my device, and my location.

This issue is not new; the same issue of civil location has arisen in the transportation industry. In an aircraft on an international flight, the *Tokyo Convention* specifies that with an aircraft in flight, a passenger is "located" in the jurisdiction of the country of registration of the aircraft, irrespective of the precise location of the aircraft. In ships at sea, the jurisdiction of the country of registration of the vessel is used only when the vessel is physically located in international waters.

A Perspective from a Geolocation Provider — Maxmind

A commonly used tool for IP geolocation is the *Maxmind* data. Maxmind publishes a geolocation data set that maps IP addresses into countries, and a more detailed data set that provides a similar mapping to a state and city level. The approach Maxmind uses is to gather a variety of sources about the location of an IP address, starting with the data set provided by the RIRs, geofeeds, the *Peering Data Base*, and ancillary data that can assist in triangulation of location.

The RIRs publish a report[24] relating to the allocation and assignment of IP addresses to network providers and other intermediaries. The registries contain the names of these organisations and their addresses, including a country code. This address is not necessarily the one where the addresses will be used, nor is the size of the address block the size of the announced prefix in the routing system. It's the holding organisation's address. However, in the absence of any other signal, this country code in the RIR data is as good a starting point as any to assemble a geolocation map.

From this starting point it's possible to refine this view. If a network publishes a geofeed, then this signal is useful for a specific set of address prefixes. The routing system can also be used, in that if the location where an *Autonomous System number* (AS) is being used is known, then that infers that the location of all the IP addresses originated by this AS can also be derived, in that their location is bounded by the AS's usage bounds. There's other third-party bulk data, particularly relating to areas of the world where the location service provider is not well resourced to be able to perform the data gathering and analysis directly. Also *traceroute* and *ping* data can assist in the triangulation exercise, in addition to reverse *Domain Name System* (DNS) names, which often incorporate geonames such as airport codes, particularly when naming infrastructure elements. Individually notified amendment requests, where more specific location information is submitted to the location service operator, can also be helpful. This data is fed into a set of heuristics and preference rules to create confidence assessments and conflict resolution that is used to determine a location result for all routed IP addresses.

Part of the challenge of this work is that IP addresses don't necessarily have a simple single view of the location where the address is used. There is (sometimes subtle) distinction between the location of the individual using the device, the location of the device, and the location where the traffic enters and leaves the public Internet. This distinction is significant when the device is connected by a VPN or a relay, or when the user is using a remote access tool. Maxmind's priority is to geolocate the person who is using the device.

Maxmind is also apparently looking at ways to expose their confidence intervals relating to location of observed traffic. Generally, observed network traffic for a prefix lies in a single locale, although there are signals of traffic use across larger locales reflecting mobile use and business and residential use of VPNs. The aim here with the exposed confidence level is to provide greater transparency of the calculations that result in a single-line "best estimate" location.

What about the consumers of Maxmind's data and location data in general? When you ask "where is this IP address located?", there is a *regulatory location* that attempts to identify which regulations or content access constraints (*Digital Rights Management*)[25] must be applied to an individual user at a physical location, as compared to a *network location* that attempts to make a decision between potential servers to optimise service performance between server and user device for the transaction that is optimised against this network location. Compounding this process is the desire to perform both regulatory and network location at once from a single location signal.

Consumers of this location data tend to have a wide diversity of requirements from location information. The Digital Rights Management and regulatory compliance activities tend to seek the most precise location possible to obtain, and to do so with the most recent data that is possible to obtain in order to avoid misapplication of these location-based constraints.

Ad placement technology is concerned more with demographics and coverage. High precision location is relevant only to the extent that it can inform a demographic indicator (such as working in a particular sector or being on a university campus). There is also the security perspective, where location data can be used to identify anomalous transactions that are potentially fraudulent.

In the case of IPv4, and also increasingly in IPv6, there is the issue of shared addresses, where a single public address may be shared among multiple users, either serially or simultaneously. There is no clear answer here as to how to resolve this problem, and the consumer of IP location data should be aware that there is not necessarily a tight correlation between a single individual user (at a single physical location) and a single IP address, and this reality will necessarily add some variability to the resultant location information for some IP addresses. For example, in a mobile context the radius of uncertainty of a location given in latitude/longitude coordinates may be many hundreds of kilometres because the address may be observed to be in use across all of these potential areas. On the other hand, too precise a location can become personally identifying information and potentially cause harm to users. In Maxmind's case they limit the precision of a radius of uncertainly to no smaller than 5km, because of concerns related to the protection of individual privacy if information about a more precise location is used in this context.

However, the role of geofeed data is important in this area of associating an IP address with a location, and some sites provide some insights as to who uses this geofeed data. **`https://geolocatemuch.com/`** is a very good example of the location providers that use this geofeed data.

Many of the consumers of this data are motivated by a desire to comply with location-based constraints, such as digital rights management where a provider has purchased the right to retail content or service to a particular geographic territory. A related issue concerns legal compliance, including compliance with government-imposed sanctions, and the levying of taxes and tariffs on the provision of digital goods and services that may apply depending on the geolocation of the provider and consumer. Many who use this data are very interested in both access to data that reflects an up-to-the-second view of a user's location and highly specific locations.

Another community uses location as a demographic attribute to profile users. A good example is the ad placement industry, where an advertiser may want to purchase the placement of online ads that are directed to particular locations, in addition to other demographic profile specifications. The value of this data to the ad placement operator is generally lower than the value associated with legal compliance.

A further class of consumers is related to the cybersecurity activity, where the analysis of threat and the associated calculation of risk, use location as part of the inputs to such assessments.

The role of using location to meet users' expectations for content, context, and experience customization is also of interest. A user who is geolocated to Germany, for example, might expect that by default the language of presentation the browser uses would be German. Geolocation can also assist in choosing a nearby server to service a user's request, on the assumption that closer services will provide a superior user experience. Answering a user's query for a great coffee shop with the details of a shop on the other side of the world is nothing more than an exercise in artful frustration! The expectation is that the results of a query use a context of locally relevant material in preference to other material, where what is "locally relevant" is very much a question of location.

Some of the anomalies that can arise come from the misapplication of geo data. A service provider may have purchased a geofeed with the primary objective of optimising the user experience when accessing their service but apply the same geo data for the application of controls to comply with regulatory constraints. For this secondary objective the geo data may not be sufficiently precise, nor timely.

A perspective from a Content Distribution Network Operator — Akamai

Akamai sells a product called *EdgeScape*, which is a collection of IP geolocation feeds, using the RFC 8805 format. Like other IP geolocation providers, Akamai takes a variety of sources of IP-to-location and creates a single outcome based on confidence scoring of the various inputs. They also use the same data for their internal use, such as a selective service to optimize the responsiveness of service requests. (Figure 1)

Figure 1: IP:Geo Lifecycle at Akamai.

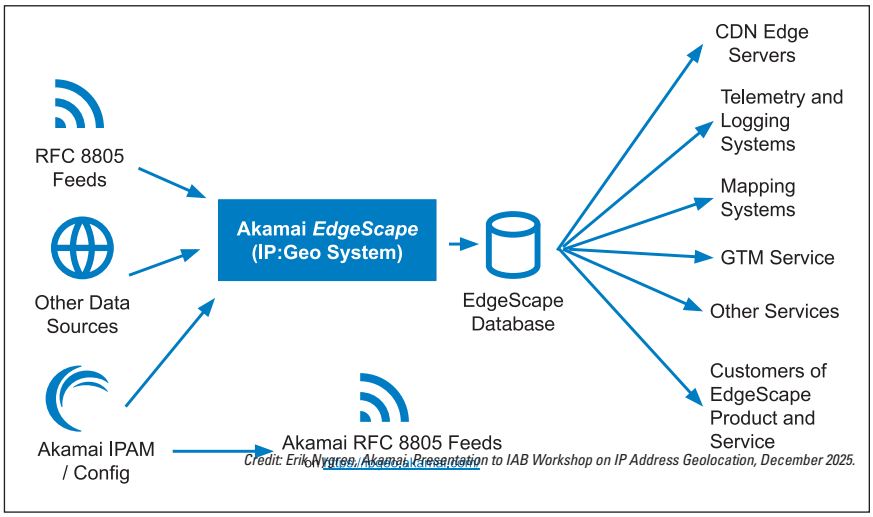

What Akamai produces is a deliberately coarse approximation of the location as a publicly routable IP endpoint, whether it's a client device, a DNS resolver, a data centre server, or a *Network Address Translation* (NAT) exit point of a geolocation sentinel.

As a cloud content provider, they use a location-informed DNS server to direct a user to a nearby (in network terms) content server. The distinction here between this form of IP geolocation and that generated by Maxmind is that it is the network gateway that's important to Akamai, and not necessarily the user's civil or geographic location. Thus, for VPN users, it's the VPN egress point that Akamai's geolocation service identifies. For geolocation-informed DNS, they need to perform IP location at high volume and low delay.

The geofeeds that networks publish are valuable in assembling this data, but they are not necessarily sufficient. Many network operators do not self-publish a geofeed data set. Those data sets that are published are not always accessible, and there is an unknown delay between the time of gathering the information used to inform the geofeed data and the publication of the feed itself. They also observe some level of discrepancy between IPv4 and IPv6 location data, where the use of the same endpoint identity token shows that this is the same endpoint device, yet the geofeed IPv4 and IPv6 data show divergence in the described locations.

A related issue is IP *reputation*, which is similar in many ways and different in some important ways. IP reputation has a very important time axis, where compromised endpoints rapidly affect the associated *reputation score*, while location changes at a far slower rate, if at all.

The RFC 8805 format does not offer a time, geolocation precision, or confidence of the data contained in the data set. It would also be useful to classify the address by network type, such as data centre or cloud infrastructure, mobile network, fixed access, or enterprise, but that is asking too much from an informal system of provider-published data.

The use by location-informed DNS is interesting, as in this case it is an instance of a third-party triangulation problem of *relative location*: How can an observer (the DNS server) determine the "best" server for a given user when the observer is neither the user nor the server? Figure 2 shows this scenario.

Figure 2: DNS Content Steering.

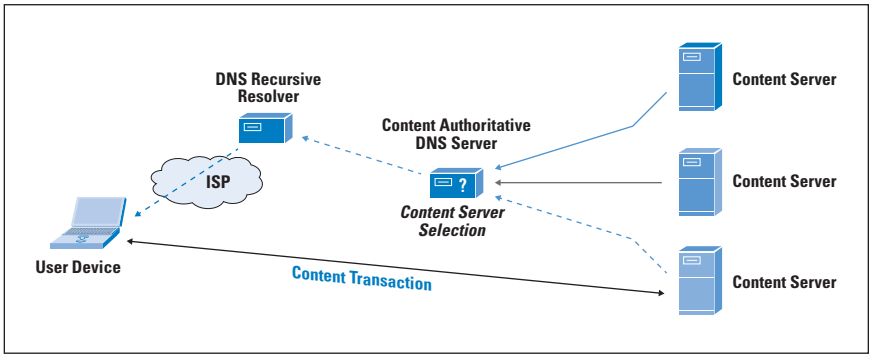

This triangulation problem also has compliance aspects; the selected server choice may be constrained by regulatory considerations as well.

The term "geolocation" is shorthand for a set of quite distinct questions. For example, when applying geolocation to a hypothetical ambulance dispatch function, these questions need to be addressed: where to send the ambulance (the user's location), which ambulance station should dispatch the ambulance (a server location), and what equipment should be in the ambulance that is dispatched (a user experience optimisation consideration)? And another question arises: what is actually needed, an ambulance or a fire truck! The "interesting" outcomes happen when you think you are asking one type of geolocation question and you get an answer to an entirely different question!

A Perspective from a Content Intermediary — Cloudflare

Cloudflare's role in web-based content and service delivery can be considered as an active intermediary between the user and the web service that the user wishes to access. The behaviour of Cloudflare's intervention is determined by the web provider for the service, not necessarily by Cloudflare. Cloudflare provides control mechanisms to web service providers, but how the web service provider sets these control mechanisms is up to them. Cloudflare's service could be described as a "content conduit." Cloudflare is also an access provider with various VPN-related access products, predominately for enterprise markets. Cloudflare consumes IP geolocation products that others provide and uses these products to map users and services to locations, but in this case Cloudflare is neither the user nor the service provider.

Cloudflare uses anycast extensively to assist in steering users to the closest service, where "close" is defined in network terms. This use implies that no infrastructure addresses used in Cloudflare's anycast *Content Distribution Network* (CDN) service have a single location, and it does not necessarily fit in a conventional model of IP geolocation that wants to associate an IP address prefix with a single location or single jurisdiction.

The gaps in geolocation that Cloudflare is aware of include:

- A lack of recorded provenance of the location mapping data.
- A lack of timings between data updates and pickup by various location data brokers.
- A lack of transparency on the data-broker decisions between potentially competing location claims.
- A lack of a framework to describe data quality and a clear way of measuring this quality.

If the ultimate objective is to map a human user to a physical location, then the use of geolocation of IP addresses is a convenient and potentially low-cost method to achieve this result, but it has some fundamental weaknesses when you consider operational practices of dynamic address assignment, address sharing through NATs, the use of active proxies as intermediaries, the variable quality of location, the opacity of operators' address plans, and others.

The technique of using IP addresses to determine geolocation is cheap, fast, and efficient, but it has some rough edges relating to determining the quality of the result.

IP addresses are not the only location proxy, and other signals are feasible in certain contexts. Web location, as provided in the *World Wide Web Consortium* (W3C) *Geolocation API*[26] can be accessed, with the user's permission, or mobile device location available at the application level is used. This location technique uses the device's GPS or mobile/Wi-Fi signals to detect location, and it operates with high precision down to the level of a small number of metres of uncertainty. This approach requires the explicit permission of the user through a browser query or configuration session, and it is commonly used when the precise location of a user is required, for example, locating a user on a map to find travel distances to a destination, or delivery of services to a residence. The high level of precision has privacy issues, and there are no explicit limitations on what the service provider can do with the data obtained in this manner, as presumably the service provider now knows the device IP address as well as its location to a high precision and could presumably sell this collected data to others.

The IP address by itself is a poor proxy for a user's location. On the other hand, if an IP address can be mapped to a location and purpose with some level of surety, then it's possible to build a reputation for that address including a profile of a "normal" pattern of network use. The address might be used by a residential access service in a fixed location or shared across residential users in a bounded locale. Or it could be used in a mobile access network, with associated broader parameters of the bounds of a locale of use. The address might be a VPN or tunnel interface, in which case the IP address is opaque as to the location of the user at the other end of the tunnel. The address might be used for network or service infrastructure, in which case it would not normally be used to initiate outbound connections. Additional metadata about the use or context of IP addresses could assist in the detection and response to unwanted or hostile traffic by network and service operators.

On the other hand, there is an obvious tension between user privacy and the potential to exploit location metadata to harden networks to actively resist various forms of hostile attack. Location establishment processes, authenticity of the data, and an understanding of the meaning and context of the data vary widely among data publishers, data processors, and brokers and users. It would be far easier to support users who are impacted by location attribution errors if we knew whom to talk to, how to submit corrections, and how to make available the provenance and confidence in the data.

The result we have today with IP geolocation is highly chaotic, and places probably too much reliance on inference and guesswork.

The Perspective of a Regional Internet Registry — the RIPE NCC

A different aspect of geolocation is that of identifying the location of network infrastructure, as distinct from the location of the network users.

One motivating factor is an effort to relate a natural event at a given location to assess the likely impact on network infrastructure. Another motivation is to expose the geographic path taken by Internet packets when using *traceroute*, such as **`https://geotraceroute.com/`**. This service can expose potential performance-related behaviours such as hair-pinning, when traffic passed between geographically adjacent points is passed through a remote waypoint. In addition, some network path questions are of national interest: Is traffic to/from some remote service point passed across network infrastructure located in a particular country or countries? Or does traffic between two endpoints within one country pass through other countries? Or is critical infrastructure for my country's users, such as authoritative DNS servers, for example, located in other countries? And there is the more general question of assessing the extent of Internet self-resilience for a country. To what extent are domestic Internet services reliant on services located in, and controlled by, foreign entities? All these questions necessarily involve establishing the geopolitical location of network infrastructure.

The *Réseaux IP Européens Network Coordination Centre* (RIPE NCC) has commenced an initial service to maintain a public data set that contains the geolocation information for Internet infrastructure, located by the IP address of an infrastructure element. The sources they use include the IP address databases maintained by the RIRs, which serve as a broad hint, but are of limited use in terms of precision. In addition, the Peering Database data set has reasonable accuracy, but limited coverage. There are also geofeeds, and it is also possible to mine the DNS names of infrastructure elements and then make some inferences based on the labels used in the DNS Name (such as the *The Aleph Project*[27] at Northwestern University in the United States). Another option is to use crowdsourcing or association hints from previously positioned IP addresses to infer the location of the covering address prefix. Or basic triangulation, where a low delay between path-adjacent IP addresses can infer some level of geographic proximity. These various location hints can be assessed with a level of confidence. This tool, **`https://ipmap.ripe.net/`**, is currently under development at the RIPE NCC.

This activity, however, is not without aspects of controversy. Infrastructure maps, such as **`https://openinframap.org/`**, necessarily need careful balance between providing public information and overexposing critical service infrastructure to the potential for hostile disruption, and as we've seen in the submarine cable world in recent years, deliberate disruption of communication systems can have extremely far-reaching consequences.

Location and Regulatory Compliance — a Perspective from GeoComply

One entity has based its business model on the observation that IP geolocation is not a good fit for various overlay business models that demand high precision. For example, in the domain of sports betting in the United States, some states have compliance constraints based on tribal land boundaries, and these constraints require a level of positional precision beyond what is normally available in GeoIP. VPNs can effectively "teleport" users, and they can be used to bypass territorial business models.

The content industry has used regional pricing models for many years (for example, Amazon Prime Video is 6 times more expensive in the United States than in India, and YouTube Premium is close to 4 times more expensive in Switzerland than in Serbia), and VPNs can undermine such differential business models by teleporting a user to a region where the same service is cheaper to access.

VPN use in the retail access space is widespread, with some 1.75 billion users (one third of the total user base) regularly using VPNs for reasons that include privacy, security, bypassing local censorship actions, and access to geo-restricted content. The global VPN market is currently valued at some USD 44.6 billion, with projections for continued strong growth in the coming years. Given the deliberate manipulation of the visible IP address to use the IP address of the VPN egress point, and not that of the device, then if the objective of GeoIP is to locate the user, this method does not work very well!

In response, many parts of the industry have turned to application-based location sharing, where the platform gathers the user's consent to share the device's geolocation with the application service provider. But can an application service provider trust this location data even if the user has provided explicit consent? Not really, and as much as there are applications that can tap into a device's location there are also applications that rewrite a device's location to deliberately mislead these location-aware apps! The download statistics for GPS spoofing applications number in the 10's of millions and rising!

Location determination has an asymmetry of relative value. It may well be that the privacy of my location is important to me and I might be prepared to spend some money to ensure its privacy or obscure any location I may provide from my device, but to some providers I am just one consumer out of millions, and the value of knowledge of my location to the provider is nowhere near as valuable as concealing or obfuscating this information is to me. In many scenarios a provider is not prepared to spend the same amount of time to find out my location as I might be prepared to spend to obscure it, should I choose to do so. This picture changes to some extent when regulatory mandates may also entail large fines for breaches, increasing the motivation of the provider, but as an economic issue it's generally the case that information about my location is far more valuable to me than it may be to any third party.

The effort to establish the true location of a device necessarily needs to pull in more data from the device, including the local time of day and the nearby Wi-Fi *Service Set Identifier* (SSID) names and their signal strengths, in addition to the IP location and the reported GPS coordinates, if the device is GPS-capable and receiving a signal.

Two-factor authentication also can work in favour of geolocation, where the application requires the use of a mobile device to authenticate the access attempt and thereby also provide the device's GPS coordinates as part of the authentication process. To an extent this process relies on user consent and testing whether the device is deliberately obfuscating its location. On the other hand, in terms of identifying and countering various forms of fraud and deception, establishing a user's usual pattern of locations allows non-local access attempts to be readily identified and flagged, adding a further measure of protection for the user.

The Perspective of a Satellite Operator — Starlink

Starlink is affecting the Internet significantly. Starlink is a service based on *Low Earth Orbit* (LEO) satellites[28]. The proximity of the spacecraft to the earth-based user means that it is possible to provide a high-speed, low-latency service that is competitive with existing terrestrial-based Internet access networks. It has a further feature that allows spacecraft to relay IP packets between each other at high speed (currently 100G), which allows a user to be arbitrarily distant from the ground station that they are using. Starlink services are accessible across the entire surface of the earth, including on ships at sea in international waters and in aircraft in flight. (Figure 3)

Figure 3: Satellites Map as of December 2025— **`https://satellitemap.space`**

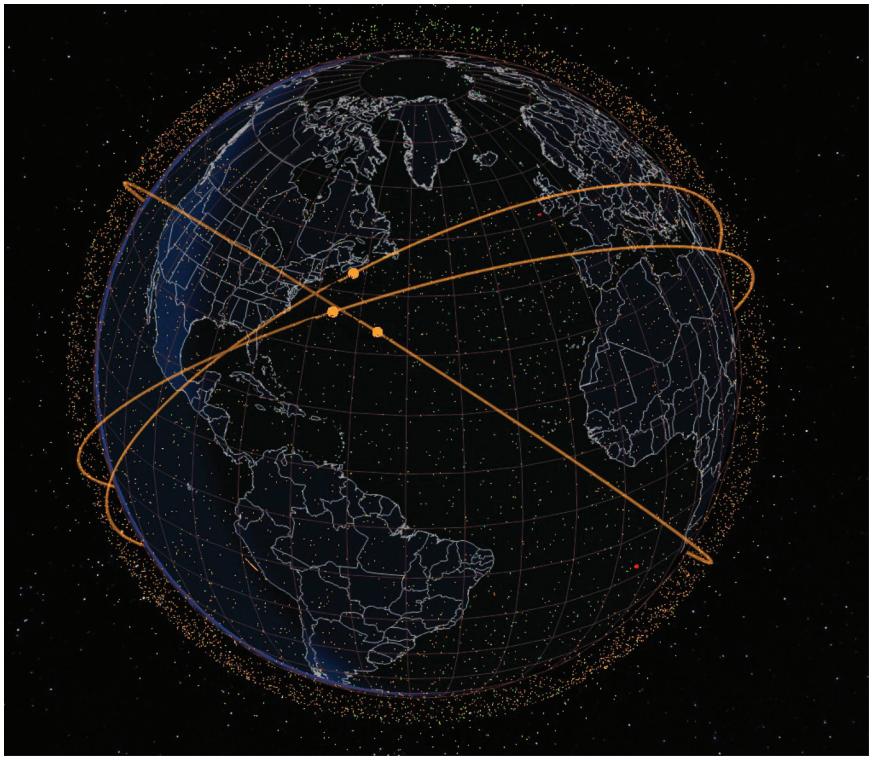

From Starlink's perspective, they want to route their users to a ground station that is the lowest latency point that can provide region-appropriate content. From a Starlink user's perspective, they expect their experience as a Starlink user to be equivalent to that of a user in a similar location using a terrestrial access service.

The question changes somewhat when the user is in international waters or in extraterritorial space such as the polar ice caps (for as long as they still exist!). Where should a geolocation service "map" users at such locations? The question of location also has implications for mobility. With Starlink, an aircraft in flight will retain the same public IP addresses for the duration of the flight, even though the take-off and landing points may be half a world apart. The compromise Starlink has adopted is to select a fixed location mapping in terms of reported geolocation country and region, even if the user (or aircraft) is roaming.

In common with various forms of tunnels and relays, the user's location and the locations where the Starlink network interfaces, *Points of Presences* (PoPs) with the public Internet are not the same (for example, many French Starlink users appear to be using Starlink PoPs located in Germany, Italy, or the United Kingdom). Should such a distinction be reflected in geofeed data? It appears to be feasible to extend the geofeed specification to convey a PoP location as well as the user's location. The data could also include an indication of whether the endpoint is roaming, although that introduces a time element that is well outside of the current scope of the geofeed data.

Starlink's preferred outcome is to allow the network to optimise traffic based on network topology rather than physical geography and embed user location signals into the application environment rather than rely on an inference drawn from the IP address. (A good example of this abstraction of location further up the protocol stack is the use of location metadata in DNS queries that include *Explicit Client Subnet* information in the DNS query.)

In spite of this apparent desire to shift away from an IP address location signal, Starlink publishes a geofeed, **`https://geoip.starlinkisp.net/`**, that maps IP addresses to countries, regions, and cities.

It is not the location of the ground station PoPs, **`https://geoip.starlinkisp.net/pops.csv`**, that interfaces to the Internet, and not necessarily the location of the Starlink terminal equipment, as evidenced by Starlink Roaming services. It appears that on the whole the Starlink geofeed locates the user to the location provided to Starlink when this user executed a service contract with Starlink. An *administrative location* seems to be an appropriate description of this particular geofeed published by Starlink.

More LEO-based satellite services are coming. The preferred outcome would be that these LEO satellite constellation operators could work with the standards developers to devise a common reference model to describe location, and a common reporting framework for location mapping.

Starlink services to aircraft reopen a topic about the location of aircraft in flight. In terms of legal jurisdiction, the *Tokyo Convention* says that the people on the aircraft are subject to the laws of the country of registration of the aircraft when the aircraft is in flight on an international route. Should the IP addresses obtained by the aircraft's on-board Wi-Fi service reflect this? What happens when an airline runs a service between two ports that are in the same country, but not that of the airline, such as a Qantas flight operated between Los Angeles and New York? Is the answer the same if the aircraft on such a flight overflies into international waters or other countries? If the answer indicates that a change in the jurisdictional compliance location is appropriate, are geofeed systems capable of reacting to this update in notional location of an aircraft's passengers in anything like real time?

Current systems have a lag that is measured in days, not minutes. Even a moving event such as the IETF meetings, have problems in this space when the event moves locations just three times a year.

Reflections — CDNs and IP Geolocation

One of the goals of geolocation for Content Distribution Networks is to associate the user with the optimal selection of an edge cache, so that the content can be delivered reliably and at the desired speed.

Much of the task can be passed to the DNS, where the CDN assumes that the user and the DNS recursive resolver that handles the user's queries are co-located to some acceptable level of proximity, and the processing of a DNS request relating to a service name that is served by a CDN platform is based on a triangulation of returning the IP address of the cache server that is considered to be closest to the client's DNS resolver (where "close" is a network-based distance as distinct from a geographic distance).

When a CDN attempts to make a selection that is optimal for the user, there are considerations of client/server latency and network path capacity, and potential backup selections to be used when the session is disrupted. What is termed "hyperlocal" data, such as local weather content, is also involved, in addition to considerations of copyright and licensing and geoblocking that can remove some candidate server selections. For many CDNs it is feasible to maintain a *proximity and performance* map, as distinct from a geographic map, and make server selections based on recently recorded performance data as well as the proximity of the end user to the candidate edge caches. This map also has an internal counterpart, where these edge caches are fed by regional caches and sent back to parent caches and the original origin source or sources (Figure 4).

Figure 4: CDNs Choosing the Best Path Exclusively using Proximity.

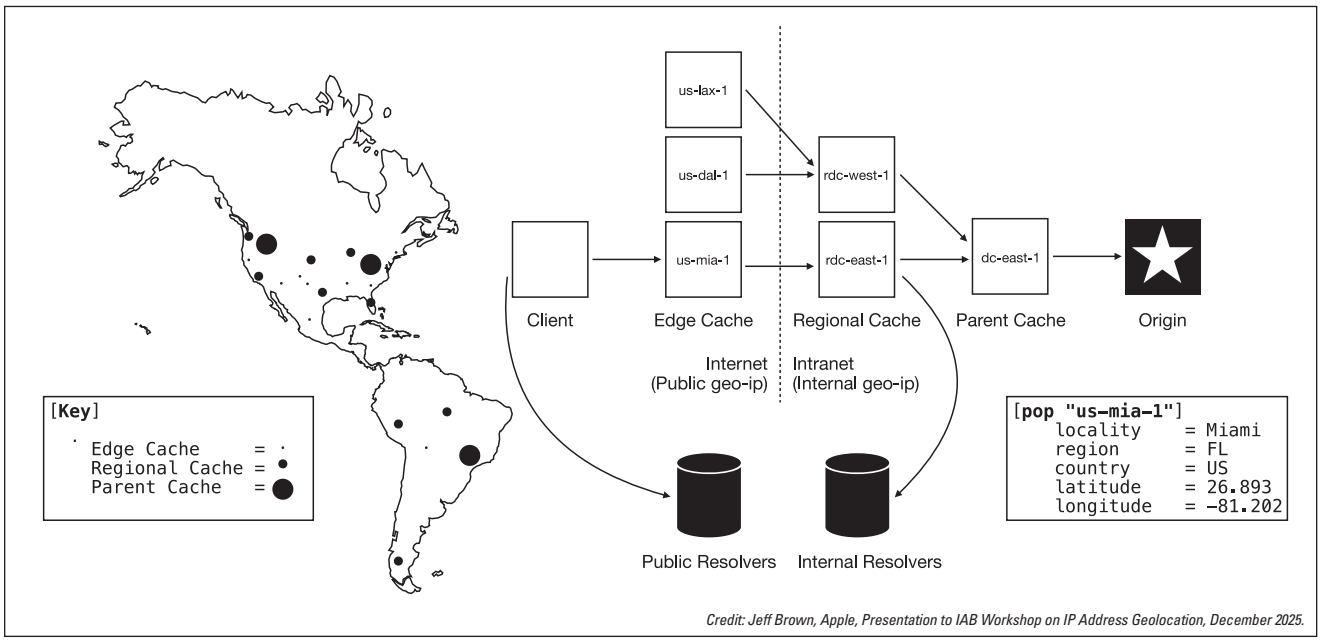

Credit: Jeff Brown, Apple, Presentation to IAB Workshop on IP Address Geolocation, December 2025.

It's also possible to use a *connection race*, where a service request triggers multiple servers (and potentially uses both IPv4 and IPv6) to send the initial transport handshake connection, asks the client to select the first response it received, and resets all the other attempts. For long-held content sessions it is also possible to *dither* the local cache selection and periodically have alternate local caches send content chunks, and if the performance of the alternate cache server is superior, then the session can shift over to the other cache server.

A slightly different approach would be to increase the amount of information provided in the DNS response to inform a client-based decision process to select the optimal server. This function can also be implemented as a server decision using the HTTP protocol, where the client can send their network geolocation to an HTTP server via an HTTP header field. This approach can provide geolocation accuracy to a level of precision determined by the client, and it gives clients more transparency regarding to whom they are disclosing their geolocation coordinates. This HTTP interaction can be extremely useful in cases where NATs, relays, and VPN tunnels are used (negating the ability to geolocate based on IP address), assuming the client is willing to disclose their location to a server. Another alternative approach is possible when using *Obfuscated HTTP* (OHTTP), where the OHTTP is aware of the client's IP address and performs a geolocation lookup of this IP address, and passes just the GeoIP record and not the IP address to the OHTTP gateway. A *geohints* draft[29] describes this latter approach.

There are multiple types of such IP teleportation proxies; for example, VPNs offer a large set of egress choices on a country-by-country basis as the core of their business differential value.

The aim of such VPN services is not to provide an accurate GeoIP location, but to do just enough to meet the distant location's compliance criteria to allow the VPN client to access content that is otherwise not permitted in their physical location.

On the other hand, simply obscuring location and other aspects of the user's profile can be challenging. In October 2025 Google shut down its *Privacy Sandbox* initiative, which was launched in 2019 to develop privacy-protecting technologies for online advertising as an alternative to third-party cookies[30]. The project aimed to balance user privacy with the needs of the ad industry, but its technologies, including the *Attribution Reporting API*, *Protected Audience*, and *Topics API*, are now being retired. Given the huge market share of Google's *Chrome* browser, efforts by Google to make it extremely easy for users to limit the extent of data generation that allows others to profile them, including location, of course, naturally attract the attention of governments and other service providers, and the outcomes of such close scrutiny are at best highly unpredictable! The compromise Google devised was a web-based mechanism that still has aspects of location reporting, but at a level where the granularity is intentionally at some half a million users or so.

IP Geofeeds

Geofeeds are an important part of the ecosystem for IP geolocation, and they continue to grow in levels of adoption (Figure 5).

Figure 5: Adoption of IP Geofeeds (from IPinfo).

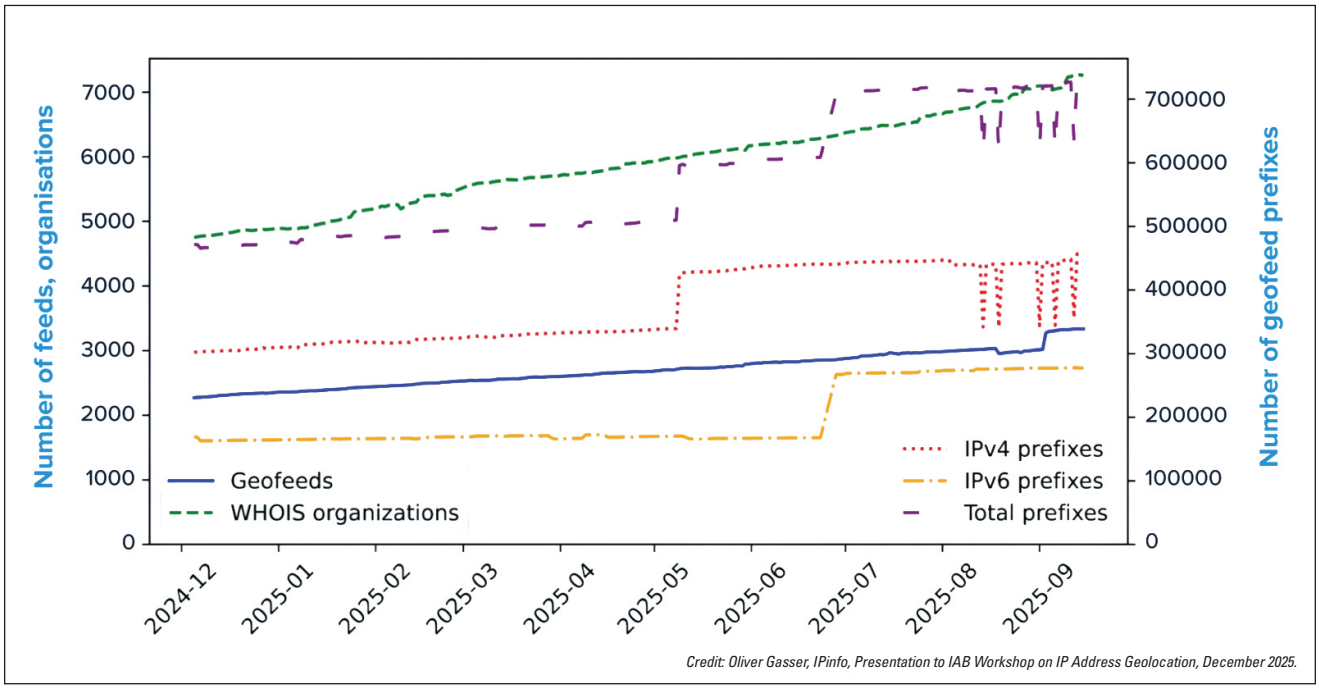

Geofeeds are one of the few ways where network operators can publish their location information in a manner that can be accessed and parsed in an interoperable manner. They are far from perfect (as can be seen from some weaknesses listed later in this section), but it's not clear if there is the common will to invest more effort into assembling and maintaining a richer geofeed data set, nor is it clear how the network operator could monetise this effort. Geofeeds are certainly popular these days, but there is no alternative out there that is comparably priced, so in a way we put up with the issues with geofeeds because that's all we have and all we can afford to do!

Research work comparing *traceroute* data with geofeed data sets at the University of California Irvine that was reported at the IAB workshop indicated the relatively poor quality of the data, where some 5% of IP addresses have anomalous location information, and one of them appears to lie in infrastructure addresses, as some 55% of traceroutes include these anomalies. (Figure 6). Experimental techniques can resolve many of these anomalies, but the question remains whether the effort required to perform this anomaly resolution results in a tangible increase in the value of the aggregated geofeed data.

Figure 6: IP Geofeed Anomalies.

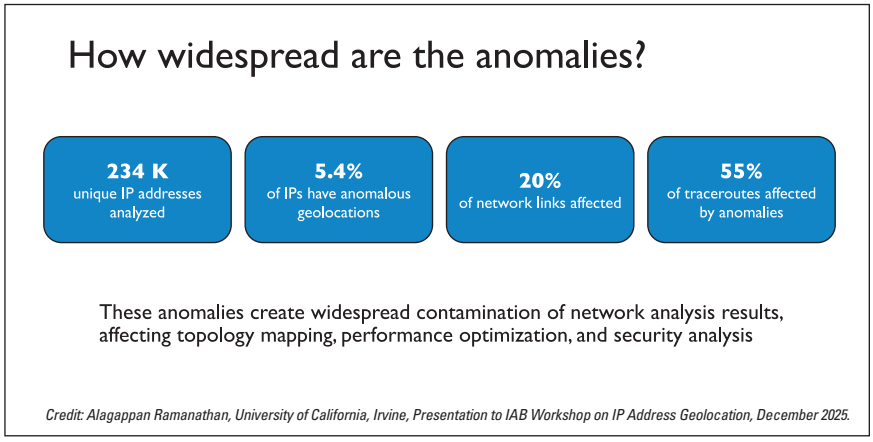

Applying *traceroute* times to geolocation questions has some weaknesses. The *traceroute* tool cannot report one-way delays, it reports only the round-trip time from the sender to the reporting point and back. If the return path varies for the intermediate points on the forward path, then *traceroute* itself reports anomalies that are challenging to analyse. In the exercise, the researchers use *Paris traceroute* and multiple *traceroutes* from multiple points to try to address any such anomalies, but it's unclear how effective it is.

What are the problems here in terms of trust and accountability? Geofeeds are a collection of self-assertions by a collection of entities who are not directly accountable to any particular public function or authority.

It seems somewhat anomalous to see various forms of government regulation and even legislation that relies on a foundation of trustable and authoritative data that assumes the existence of a working framework that can reliably locate users and the devices they are using in a country, a state, or even a local authority when in fact this foundation data is fundamentally fractured! However, public authorities, as well as industry's content distribution systems today, have locked in an assumption that geolocation is available and is working acceptably well. Imperfect as it may be, it's something for which we really have to accommodate its manifest inaccuracies and potential misdirections.

Geofeeds are certainly convenient, but they do not withstand critical scrutiny in terms of being fit for purpose. The issues with geofeeds include:

- They have an implied precision, but should that be explicitly described?
- They use civil names for higher-precision locations, but do not provide a reference set of descriptors for such civil names, nor a character set to use.
- They have no alternative location descriptors, such as latitude/longitude coordinates.
- They do not add attributes to the IP address prefix to designate unicast vs. anycast and shared NAT use vs. dedicated use vs. infrastructure use.
- Mobility is an unaddressed problem.
- They do not describe the source of the geofeed data; they rely on the location of the feed from where it was retrieved to provide an implicit signal of authenticity.
- They do not include times relating to use and currency of the data.
- They contain no implicit signal of user awareness and consent. There are many nuances for consent. How can a geofeed convey such nuances?
- They have no way to describe confidence or authenticity in the accuracy of the data.
- No mechanism is provided for distribution of this data. How to discover a geofeed and how to assess the confidence in the authenticity of a geofeed are not adequately addressed.
- How are errors noted in a geofeed, and who is the relevant contact point for amendments to a geofeed?
- Who maintains geofeed data elements? How frequently? Under what authority?

This situation is clearly not ideal.

It is definitely not ideal for the parties who have an interest in location data. A lot of the time this data and its associated framework of operation are so far from ideal that it cannot be trusted for all but the coarsest of positioning tasks. But the industry has managed to adapt to what it does today, adapting to both its weaknesses and definition of the role it plays in today's environment. You'd probably go hungry if you relied on IP geolocation to get a pizza delivered to your door! On the other hand, you would likely be able to stream country-limited content to your device, or get a reasonably relevant weather forecast from your residential connection using the same geolocation data.

Are IP Addresses the "Right" Tool for Location?

Perhaps these issues are fundamental to the approach of using IP addresses as a proxy for geolocation. Semantically, IP addresses were originally intended to convey the *identity* of an attached endpoint and the network-relative *location* of this endpoint. The intended outcome was to be able to pass a packet into a hop-by-hop stateless forwarding framework and have the packet reach the intended destination. The intended destination was identified by its interface IP address (identity), and the sequence of network forwarding decisions to pass the packet through the network were informed only by this destination address (network location).

However, network location is not necessarily physical location. IP addresses are not distributed according to a geographic hierarchy, and they have no geographic significance. IP addresses are distributed in a manner aligned with individual network providers. It is a happy coincidence that most network providers operate within a national jurisdiction, so the IP addresses that they deploy to meet their customers' needs aligns with the country in which they operate. So, to a first-order approximation with many exceptions, IP addresses can be aligned to individual countries. This alignment can be seen in the reports generated by the address distribution agencies, the RIRs, in their data reports[24]. It must be stressed that the geolocation was not a motivation behind these reports. The country code in these reports reflects the location of the head office of the enterprise that received this address allocation, not the location where these addresses are used on the Internet.

These address holding entities may have operations in many countries where they deploy these addresses. This data is not necessarily reported to the RIRs, and the RIRs do not generally publish it. The end user may use VPNs, relays, proxies, or other forms of encapsulation or obfuscation to deliberately shift their apparent location elsewhere or conceal their local IP address completely. Numerous reasons, including address scarcity and anycast services, caused many service providers to deploy address sharing technologies that deliberately break any unique association of an end device and the IP address that can be used to direct packets to it.

If an IP address can be used by many devices in many locations, then the objective to associate a physical location with each IP address is fundamentally undermined. The problem here is that an IP address does not self-identify its context of use. While we continue to use IP addresses in a manner that is still consistent with the original view of being a token with identity and network location properties, the cracks in this approach are clearly evident.

Alternatives to IP Geofeed Data

If not IP addresses, then what else is there?

There are many forms of location identification; for example, latitude/longitude coordinates, or the associated geohash framework; country names, or identified region names with a country; city names; and postal addresses that combine these sets of civil identifiers to a level of location of residences or workplaces.

In addition, means of location discovery are available, including the use of global navigation systems (GPS, Galileo, GLONASS, and BeiDou) to perform self-discovery of location; associative location discovery using Wi-Fi, *Service Set Identifiers* (SSIDs), or mobile cell tower identifiers; and the recycling of the surveying practice of digital triangulation with accurate timing and reference locations. All of these alternatives have attributes of cost, operation, precision, accessibility, spoofability, and deniability.

IP location is not useful for a pizza delivery business. But it is tolerably useful for compliance with national regulatory constraints, and the exceptions in country-level locations are adequately offset by the high availably of this data, low cost of access, and permissionless use.

If we wanted to use a different location signal, the attributes we would look for include high reliability (not necessarily high accuracy, but reliable in terms of its intended accuracy), high trust (including resistance to spoofing by the end user or by third parties), high availability, and low cost. It would also be good it we could find an acceptable compromise between privacy and accessibility that avoids both consent exhaustion and data theft.

One possible approach borrows from the Web Public Key Infrastructure (where a selection of trusted *Certification Authorities* verifies that a party has control of a domain name), where a selection of trusted location verifiers confirms the location of a user's device and issues a certificate to that effect; then the user's device can load it. The user application can then present its location certificate and some proof of possession of the device private key used in the location certificate request to attest its location to a server. The invocation of keys and certificates may give the appearance of invulnerability to subterfuge or spoofing, but that is a misleading appearance. The certificate issued by a location verifier attests that it applied some checks on certain data and is representing that the tests indicated that the device was at a certain location at a certain time.

Nothing stops the device from being physically moved to an entirely different location after location verification. Or allowing the device to export both its location keys and location certificate to a different location.

Perhaps there is more here despite the overt issues with certificate and key management. This approach changes the concept of how location is conveyed to an interested provider. Instead of having a third party generate its own map of where it believes devices are located by recording the IP address used by each device and the associated device location, have the other party to a transaction directly interrogate the client device as to its location; then it could request that the client device furnish its location and some credentials to allow this assertion of location to be verified (to some extent).

The RFC 8805 IP Geofeed location coordinates using cities is deeply unsatisfactory in terms of the variability in precision and personally identifying information. For example, a geofeed entry for "**`AU,AU-NSW,Wombat`**" places you in a pool of just 225 individuals, while "**`AU,AU-NSW,Sydney`**" identifies a pool of 5.6M individuals.

The concept of a *geohash*[31] has had some traction as a potential alternative. A geohash encodes latitude and longitude coordinates into an alphanumeric string. It is a hierarchical, grid-based system where the length of the string determines the precision, with longer strings representing smaller, more precise areas. Geohashes are useful for spatial indexing, making it faster and more efficient to perform proximity searches in location-based services by converting complex coordinate data into a single value. The selection of precision is context-sensitive. I might want to pass the pizza delivery outfit a geohash value that uniquely identifies my residence, while at the same time I might want to obfuscate the geohash value that I pass to a CDN to be at the level of a city.

The weakness of the geohash system is that it does not align to geopolitical boundaries, so its use in contexts of compliance is limited unless you select a high-precision geohash. If you make the grid size of the geohash framework large enough to encompass the entire country, in many cases it's likely that the geohash cell will include not just the intended country, but adjacent countries as well.

The next likely steps include an effort to develop some standards relating to latency-based server selection, adding *Services Binding* (SVCB) parameters relating to network location into DNS responses, and standardising IP geolocation HTTP client hints. These steps work to a common objective of putting the end client back in the loop about the disclosure and use of geolocation, which is substantially different from today's framework, which is largely driven by network and service operators without any engagement by (or even information given to) clients of any location tracking that is taking place.

What do we want from a location framework? That's a harder question, and I suspect that answers vary depending on whom you ask and what your use case might be.

- *Accuracy*. When I am using a location tracker to find my misplaced keys, accuracy to a meter or so is really helpful. On the other hand, if a social network platform wants to display my current location, then I might find an accuracy of a few tens of kilometres to be an uncomfortable intrusion on my expectations of personal privacy.
- *Privacy*. To what extent can I control the dissemination of my current location? Or my location history? Can third parties derive my location without my consent and then sell this information to others? It appears to be what happens with advertising platforms who might purchase a location service and use this information to guide advertisement placements. Privacy is not an absolute attribute. It is probably more useful to think of it as the level of correlation between expectations and reality.
- *Public or Private?* Is my location part of some public domain of data? Or is my location personal and private?
- *Control*. Can I control the accuracy of my reported location and who may access this data? Can I deliberately misrepresent my location? Is my location part of a portfolio of personal information that should remain under my control at all times?
- *Authority*. Who is asserting that this is my location? Under what authority are they making this assertion?
- *Verifiability*. Is my reported location "true?" Are the physical coordinates, with the bounds of accuracy, true? Is this really me?
- *Consistency*. Is this the location of me, my device, my residence, or my workplace? Is this a physical location, or a geopolitical location, or a location relative to network infrastructure? How current is this reported location?
- *Scalability*. Is this framework one that can encompass billions of users and potentially hundreds of billions of devices?
- *Politics and Law*. Are the locations of all the people who are in a country part of a portfolio of national critical data? Should such data be accessible, or even derivable, by foreign entities, public or private? Are there distinctions that should be made with respect to a country's citizens and so-called aliens? Is my location data protected under law? What constitutes "unauthorized use" of my location data?

While the present location infrastructure is unsatisfactory in many ways, a common understanding of what can be improved—not to mention who is willing to invest the effort in defining, seeding, and maintaining an improved location infrastructure—is clearly not evident so far. It seems that each party sees a different "solution path" that would address its needs, and any evidence that these differing objectives can be coalesced into a single path is negligible.

The next big steps appear to point to application-based location—but the problem is that no application has the momentum or value to sustain a framework custom geolocation solution per application, and the application environment is fragmented to the extent that a common application-level framework is left to the Android and Apple platforms. We see that *Two-Factor-Authentication* (2FA) has moved into that space already, and the rest are following; in other words, location is a transactional—not a constant—attribute.

Moving to gathering a user's location as a transaction at the application level does allow a more explicit form of user consent. It seems reasonable to require that applications have clear privacy policies explaining data collection and sharing and require explicit user consent (opt-in/opt-out) to sell or share location data, especially for advertising or analytics. For an app to access the device platform's location, the developer needs to justify its use, and allow users to manage these permissions in settings.

However, theory and practice often head down divergent paths. Researchers from the *International Computer Science Institute* (ICSI) found up to 1,325 Android apps that were gathering data from devices, even after people explicitly denied them permission. Serge Egelman, director of usable security and privacy research at the ICSI, presented the study in late June at the Federal Trade Commission's *PrivacyCon*: "Fundamentally, consumers have very few tools and cues that they can use to reasonably control their privacy and make decisions about it," Egelman said at the conference. "If app developers can just circumvent the system, then asking consumers for permission is relatively meaningless." [32]

Reflections on Privacy and Location

I suspect that most people would agree that an individual's precise location is a piece of personal data that falls within the general topic of personal privacy. Now from that point you could adopt the attitude that in today's world personal privacy is merely an historic concept, and we should all just get over it. Or we could take the United States stance and say that whatever someone does with your personal data, as long as they declare in advance the full scope of what they might do with your data, then it's all good. Or there is the European Union position with its *General Data Protection Regulation* (GDPR) that provides all European Union citizens with strong rights over their personal data and its use.

We see today a spectrum of applications of personal location in general: the activities of advertising and customisation of the user experience, mapping location to selection of language of service, local clock setting, and selection of advertisement content; also the aspect of security and protection, using your location to detect anomalous activity that uses your identity; and then the compliance obligations where content and service providers are constrained by law to deny you access to certain content and services if you are located within their jurisdictional boundaries.

User consent is a core principle of data protection and privacy. Use of metadata obtained from network infrastructure using IP addresses to map the user-side endpoint of a transaction to a location could be seen as a grave abuse of an individual user's privacy, particularly when the harvesting and analysis of this data occurs without the users' knowledge or explicit informed consent.

Government actions tend to show a high degree of inconsistency on this space. Many countries enact regulations that prohibit access to certain classes of content and services. Governments can enact such measures only within their own territory, so any service provider who wants to comply with such measures is obliged to gather location information for a user in order to determine which content constraints are applicable for that individual user. At the same time, in many countries the collection and use of personal identification information without the user's informed consent contravenes privacy laws, and may place the service provider in legal jeopardy.

The use of IP addresses for geolocation fall into a very grey zone from this respect. An IP address that is statically and exclusively assigned to a residence may well be considered a personally identifying data point and should be treated as such. A public IP address that is used by a large-scale NAT and is shared across hundreds of thousands of users has almost no residual capability to identify any individual user. Can you tell the difference between these two forms of IP address use when looking at the IP addresses used in an individual transaction? Obviously not.

What are a user's expectations regarding privacy? In general, most free and democratic societies seek to protect a biographical core of personal information that individuals wish to maintain and control from dissemination to the state. The biographical core is not limited to identity but includes information that tends to reveal intimate details of the lifestyle and personal choices of the individual. A recent (March 2024) Judgment of the Supreme Court of Canada on this topic is well worth reading[33].

It's not just public authorities. Business obligations can be at odds with informed consent as well. Content licensing agreements require coercion of users to have their location revealed, and the business arrangements are not viable otherwise. The outcome is that the logic of business interests apparently overrides the interests of individuals when infringing on their expectations of privacy, and it leaves the individual user with minimal, if any, privacy protections under the law.

This line of reasoning can lead to a conclusion that locating a user should be handled at a level of the use of individual applications, and informed consent should be obtained by the application. In some cases this consent is already obtained, such as the controls for *Location Services* in the iPhone environment, and the W3C *Geolocation API*.

(Even so, the risks of "consent exhaustion" should be factored in, where the constant stream of popups demanding a renewal of consent to use location information does little more than motivate users to simply click through to resume what they were trying to achieve in the first place.) The current status quo of the unconstrained dissemination of geofeed data that allows any party to map IP addresses to location without informing the user of this activity—let alone the user's consent—should be at a very minimum discouraged. We don't even have a common understanding of what is tolerable as a compromise between geographical precision and respect for personal privacy.

It may well be that IP addresses are simply the wrong starting place to fulfil these desires relating to compliance, security, customisation, and performance: "You cannot get to where you want to go to from where you appear to be!"

Conclusion: Time to Replace IP Geolocation

IP geolocation is an outdated technology that creates conflicts between utility, legal necessity, and user privacy. Better solutions are feasible, but the community needs to implement and standardize. Moving away from IP geolocation benefits user privacy, web services, and network operators.

Disclaimer

The views expressed in this article do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Materials

The materials relating to the IAB Workshop on IP Address Geolocation can be found at **https://datatracker.ietf.org/group/ipgeows/about/**

References

[1] *Wikipedia* article on Satellite navigation:
https://en.wikipedia.org/wiki/Satellite_navigation

[2] *Wikipedia* article on Real-time kinematic positioning:
https://en.wikipedia.org/wiki/Real-time_kinematic_positioning

[3] Australian Government, Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts, "Social media minimum age":
https://www.infrastructure.gov.au/media-communications/internet/online-safety/social-media-minimum-age

[4] World Wide Web Consortium, "IAB/W3C Workshop on Age-Based Restrictions on Content Access,"
https://www.w3.org/events/workshops/2025/iab-w3c-workshop-on-age-based-restrictions-on-content-access/

[5] *Wikipedia* article on Global Positioning System:
https://en.wikipedia.org/wiki/Global_Positioning_System

[6] IETF Datatracker, "IAB Workshop on IP Address Geolocation," **https://datatracker.ietf.org/group/ipgeows/about/**

[7] *Wikipedia* article on E.164: **https://en.wikipedia.org/wiki/E.164**

[8] *Wikipedia* article on Telephone Exchanges: **https://en.wikipedia.org/wiki/Telephone_exchange**

[9] Daniel Karrenberg, Gerard Ross, Paul Wilson, and Leslie Nobile, "Development of the Regional Internet Registry System," *The Internet Protocol Journal*, Volume 4, No. 4, December 2001.

[10] Erik Kline, Krzysztof Duleba, Zoltan Szamonek, Stefan Moser, and Warren Kumari, "A Format for Self-Published IP Geolocation Feeds," RFC 8805, August 2020.

[11] International Organization for Standardization, "ISO 3166 Country Codes." **https://www.iso.org/iso-3166-country-codes.html**

[12] IP2Location, "IP2Location™ ISO 3166-2 Subdivision Code." **https://www.ip2location.com/free/iso3166-2**

[13] Randy Bush, Massimo Candela, Warren Kumari, and Russ Housley, "Finding and Using Geofeed Data," RFC 9632, August 2024.

[14] Cengiz Alaettinoglu, Curtis Villamizar, Elise Gerich, David Kessens, David Meyer, Tony Bates, Daniel Karrenberg, and Marten Terpstra, "Routing Policy Specification Language (RPSL),"RFC 2622, June 1999.

[15] Matt Lepinski and Stephen Kent, "An Infrastructure to Support Secure Internet Routing," RFC 6480, February 2012.

[16] IETF geopriv Working Group: **https://datatracker.ietf.org/wg/geopriv/about/**

[17] Richard Barnes, Matt Lepinski, Alissa Cooper, John Morris, Hannes Tschofenig, and Henning Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications," RFC 6280, July 2011.

[18] Alexander Mayrhofer and Christian Spanring "A Uniform Resource Identifier for Geographic Locations ('geo' URI)," RFC 5870, June 2010.

[19] Mary Barnes (editor), "HTTP-Enabled Location Delivery (HELD)," RFC 5985, September 2010.

[20] Paul Ferguson and Geoff Huston, "What is a VPN? Part I," *The Internet Protocol Journal*, Volume 1, No. 1, June 1998.

[21] Paul Ferguson and Geoff Huston, "What is a VPN? Part II," *The Internet Protocol Journal*, Volume 1, No. 2, September 1998.

[22] Apple, "About iCloud Private Relay,"
https://support.apple.com/en-us/102602

[23] Maxmind: **https://www.maxmind.com/en/home**

[24] Allocation and Assignment of IP Addresses:
https://www.potaroo.net/bgp/stats/nro/delegated-extended-latest

[25] *Wikipedia* article on Digital Rights Management:
https://en.wikipedia.org/wiki/Digital_rights_management

[26] World Wide Web Consortium, "Geolocation API Specification," May 2015.
https://www.w3.org/2008/geolocation/PER-geolocation-API/

[27] The Aleph Project: **https://thealeph.ai/**

[28] Dan York and Geoff Huston, "Low Earth Orbit Satellite Systems for Internet Access," *The Internet Protocol Journal*, Volume 26, No. 2, September 2023.

[29] Tommy Pauly, David Schinazi, Ciara McMullin, and Dustin Mitchell, "The IP Geolocation HTTP Client Hint," Internet-Draft, Work in Progress, **draft-pauly-httpbis-geohash-hint**, April 2026.

[30] Arvind Hickman, "Google Shuts Down Its Privacy Sandbox," *B&T*, October 21, 2025.
https://www.bandt.com.au/google-shuts-down-its-privacy-sandbox/

[31] *Wikipedia* article on geohash:
https://en.wikipedia.org/wiki/Geohash

[32] Alfred Ng, "More than 1,000 Android apps harvest data even after you deny permissions," *CNET*, July 8, 2019.

[33] Supreme Court of Canada, "R. v. Bykovets":
https://www.scc-csc.ca/judgments-jugements/cb/2024/40269/

GEOFF HUSTON AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for building the Internet within the Australian academic and research sector in the early 1990s. He is author of numerous Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005. He served on the Board of Trustees of the Internet Society from 1992 until 2001. At various times Geoff has worked as an Internet researcher, an ISP systems architect, and a network operator. E-mail: **gih@apnic.net**

Book Review: The NIST 2.0 Cybersecurity Framework

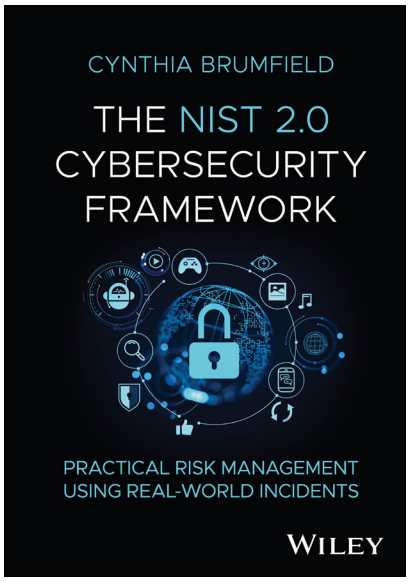

The NIST 2.0 Cybersecurity Framework: Practical Risk Management using Real-World Incidents, by Cynthia Brumfield, ISBN-13: 978-1-394-35218-0, Wiley, April 2026.

Cybersecurity expert and noted tech journalist Cynthia Brumfield has updated her seminal reference work explaining the *National Institute of Standards and Technology* (NIST) *Cybersecurity Framework*, which we'll call NIST 2.0 for brevity's sake. The book has been extensively revised for the latest version of the framework so that it can still be a potent source of information on the standards. But before I tell you why you would want the new work, let's first describe what the book is not.

First and foremost, this book is not some potboiler, wide-eyed romp through the major cybersecurity blunders of the past. In fact, Brumfield readily admitted to me in a recent conversation why her book "isn't a narrative-driven popular book, nor a technical deep-dive for engineers." It also doesn't offer a series of prescriptive security "recipes" to improve your enterprise security, or fly above in the theoretical clouds.

Instead, she tries to fill the gap between what NIST 2.0 is and how to move in its direction through better and actionable risk-management practices. It does this by using those real-world attacks and vulnerabilities as teachable moments that can demonstrate the consequences of not following the NIST dictates. Consider this more of a reference guide that will take apart the framework into its six structural components, and offer ways that practitioners can decide how to implement the framework's best security practices.

These components are divided into NIST 2.0's six categories: *identity, protection, detection, response, recovery*, and *governance*. A chapter is devoted to each category, picking apart the various components, such as asset management, inventories of hardware and software, data flows, services, and other items that make up the identity category. Each component is further dissected into steps a typical IT person would go through to improve security, various implementation examples (such as the need to integrate cybersecurity considerations into digital product life cycles), and references to various federal and other cybersecurity standards to further guide your work. At the end of each chapter is a short quiz to see if you have been paying attention, along with an extensive bibliography of news and analysis articles related to the particular category topic.

Why is this level of detail necessary, and even desirable? Mainly because cybersecurity frameworks are not very usable. Security practitioners need some document to help put things into practice. NIST's framework is widely acknowledged as a useful structure, but it doesn't offer step-by-step implementation guidance.

And to some extent, these steps are hard to codify because "the devil is in the details," and there is no one-size-fits-all, or even one-size-fits-a-couple-of-enterprises for that matter. That is where her book can be handy, especially in filling the gap between knowing and executing effective cybersecurity. "My book attempts to close that gap by translating NIST 2.0 into actionable risk-management practices grounded in real-world incidents," she told me.

If that makes you think her book could be used as a textbook for a few different computer science courses, you would be right, and that is one reason why Wiley's academic division is publishing the revised version. Brumfield shared with me one such course offered by Metro State University in St. Paul, Minnesota, on "Risk Management and controls for medical devices," a course which is certainly needed given the historically lax security procedures of many such embedded devices. (The course currently uses the earlier edition of her book.)

While Brumfield's book probably won't be bedtime reading by most of the intended audience, it does provide a clear vision of how to best implement a holistic, coherent, and thorough cybersecurity framework to take common theory into actionable practice. And I think colleges that offer computer science classes that teach how to figure out risk management and set up solid corporate security practices should add this book to their required reading lists.

—David Strom
david@strom.com

Can you hear me now?

This Automatic Electric Model 40 "Monophone" was produced between 1939 and 1957. It has been on a World Tour to demonstrate how Voice over IP can support even old equipment like this. Most recently it operated at a conference in Meißen, Germany.

Fragments

ICANN's OCTO Publishes paper on ICANN and the IETF

The *Office of the Chief Technology Officer* (OCTO) of ICANN has recently published a paper describing the relationship between ICANN and the IETF. The introduction reads as follows:

"The Internet's infrastructure is maintained by the coordinated efforts of many. The engagement of the *Internet Corporation for Assigned Names and Numbers* (ICANN) with the *Internet Engineering Task Force* (IETF) constitutes one essential thread in that fabric. Both communities work to keep the global Internet functioning as a single, interoperable network. Although ICANN and the IETF have different mandates, their missions are deeply intertwined.

This relationship is the main theme of this paper. The IETF is fundamentally important to ICANN and has been since day one of ICANN's formation.

Part One of this report explores the relationship between ICANN and the IETF. It describes how the two organizations and their communities interact, and why the time and energy ICANN commits to the work of the IETF is warranted.

Part Two provides an update on several key *Domain Name System* (DNS)-related activities that occurred at the IETF in 2025 that participants in the ICANN ecosystem may find useful and informative.

All three pillars of the ICANN community regularly participate in the IETF. This paper, however, focuses on how staff inside ICANN Org interact with the IETF."

You can read the full document here:
https://www.icann.org/en/system/files/files/octo-043-16apr26-en.pdf

---

Check your Subscription Details!

Make sure that both your postal and e-mail addresses are up-to-date since these are the only methods by which we can contact you. If you see the words "Invalid E-mail" on your printed copy, this means that we have been unable to contact you through the e-mail address on file. If this is the case, please contact us at **ipj@protocoljournal.org** with your new information. The subscription portal is located here:
https://www.ipjsubscription.org/

Read Any Good Books Lately?

Then why not share your thoughts with the readers of IPJ? We accept reviews of new titles, as well as some of the "networking classics." In some cases, we may be able to get a publisher to send you a book for review if you don't have access to it. For more information, contact us at **ipj@protocoljournal.org**

NDSS Symposium Heads to Seoul in 2027

The *Internet Society* recently announced that the *2027 Network and Distributed System Security* (NDSS) Symposium[1] will take place in Seoul, Republic of Korea, from 22–26 March 2027. Registration for the symposium will open in October 2026.

Recognized as one of the world's top four cybersecurity research conferences[2], the NDSS Symposium brings together hundreds of top scholars, academics, and practitioners to publish cutting-edge research in network and distributed systems security.

The move to Seoul, Republic of Korea, in the Asia-Pacific region reflects the globally-connected Internet and the global cybersecurity community working to keep it secure. In 2026, the Symposium welcomed 674 participants from 33 countries, continuing a steady trend of growing international engagement. Participation from the Asia-Pacific region alone has increased by more than 60% over the past three years.

"A safer Internet depends on world-class security research and strong global collaboration. Hosting NDSS Symposium in Seoul brings one of the world's top four cybersecurity research conferences closer to a rapidly growing hub of innovation, while further strengthening the international partnerships and exchange of ideas that help advance Internet security," said Dr. Joseph Lorenzo Hall, Distinguished Technologist at the Internet Society.

According to Internet Society *Pulse*[3] and the ITU's *Global Connectivity Report 2025*[4], Asia-Pacific is one of the fastest-growing Internet regions in the world. Seoul is home to world-class infrastructure, leading universities, and a vibrant technology ecosystem. It offers an ideal setting for advancing global dialogue on Internet security. Hosted by the Internet Society, the NDSS Symposium has advanced cutting-edge cybersecurity research for over three decades.

Founded in 1992 by Internet pioneers, the Internet Society is a global non-profit organization working to ensure the Internet is for everyone. Through its community of members, special interest groups, and 130+ chapters worldwide, the organization defends and promotes Internet policies, standards, and protocols that keep the Internet open, globally connected, and secure. For more information, please visit: **internetsociety.org**.

[1] **https://www.ndss-symposium.org/ndss2027/**

[2] **http://jianying.space/conference-ranking.html**

[3] **https://pulse.internetsociety.org/en/blog/2025/05/where-the-internet-is-growing-most/**

[4] **https://www.itu.int/itu-d/reports/statistics/global-connectivity-report-2025/**

Thank You!

Publication of IPJ is made possible by organizations and individuals around the world dedicated to the design, growth, evolution, and operation of the global Internet and private networks built on the Internet Protocol. The following individuals have provided support to IPJ. You can join them by visiting **http://tinyurl.com/IPJ-donate**

Kjetil Aas
Fabrizio Accatino
Michael Achola
Martin Adkins
Melchior Aelmans
Christopher Affleck
Scott Aitken
Jacobus Akkerhuis
Antonio Cuñat Alario
William Allaire
Nicola Altan
Shane Amante
Marcelo do Amaral
Matteo D'Ambrosio
Selva Anandavel
Jens Andersson
Danish Ansari
Finn Arildsen
Tim Armstrong
Richard Artes
Michael Aschwanden
David Atkins
Jac Backus
Jaime Badua
Bent Bagger
Eric Baker
Fred Baker†
Santosh Balagopalan
William Baltas
David Bandinelli
A C Barber
Benjamin Barkin-Wilkins
Ryan Barnes
Feras Batainah
Michael Bazarewsky
Robert Beckett
David Belson
Richard Bennett
Matthew Best
Hidde Beumer
Pier Paolo Biagi
Arturo Bianchi
John Bigrow
Orvar Ari Bjarnason
Tyson Blanchard
Axel Boeger
Matt Boehmer
Keith Bogart
Mirko Bonadei
Roberto Bonalumi
Lolke Boonstra
Cente Cornelis Boot
Julie Bottorff Photography
Gerry Boudreaux
Leen de Braal
Stephen Bradley
Kevin Breit
Thomas Bridge
Ilia Bromberg
Lukasz Bromirski
Václav Brožík
Christophe Brun
Gareth Bryan
Ron Buchalski
Paul Buchanan
Stefan Buckmann
Caner Budakoglu
Darrell Budic
BugWorks
Scott Burleigh
Chad Burnham
Randy Bush
Colin Butcher
Jon Harald Bøvre
Olivier Cahagne
David Fernandez Cambronero
Antoine Camerlo
Tracy Camp
Brian Candler
Fabio Caneparo
Roberto Canonico
David Cardwell
Richard Carrara
John Cavanaugh
Lj Cemeras
Dave Chapman
Stefanos Charchalakis
Molly Cheam
Christof Chen
Pierluigi Checchi
Greg Chisholm
David Chosrova
Marcin Cieslak
Lauris Cikovskis
Brad Clark
Narelle Clark
Horst Clausen
James Cliver
Guido Coenders
Robert Collet
Marcial Colomer
Joseph Connolly
Steve Corbató
Brian Courtney
Beth and Steve Crocker
Dave Crocker
Kevin Croes
John Curran
Andrew Daggers
Leslie Daigle
Sergio Danelli
André Danthine†
Morgan Davis
Jeff Day
Nicholas Dean
Fernando Saldana Del Castillo
Rodolfo Delgado-Bueno
Julien Dhallenne
Freek Dijkstra
Geert Van Dijk
David Dillow
Richard Dodsworth
Ernesto Doelling
Michael Dolan
Eugene Doroniuk
Michael Dragone
Joshua Dreier
Lutz Drink
Aaron Dudek
Dmitriy Dudko
Andrew Dul
Joan Marc Riera Duocastella
Pedro Duque
Holger Durer
Karlheinz Dölger
Mark Eanes
Andrew Edwards
Peter Robert Egli
George Ehlers
Peter Eisses
Torbjörn Eklöv
Jacobus Gerrit Elsenaar
Y Ertur
ERNW GmbH
ESdatCo
Steve Esquivel
Jay Etchings
Mikhail Evstiounin
Babatunde Faluyi
Bill Fenner
Paul Ferguson
Ricardo Ferreira
Kent Fichtner
Ulrich N Fierz
Armin Fisslthaler
Michael Fiumano
The Flirble Organisation
Mícheál Ó Foghlú
Jean-Pierre Forcioli
Gary Ford
Susan Forney†
Christopher Forsyth
Andrew Fox
Craig Fox
Fausto Franceschini
Erik Fredriksson
Valerie Fronczak
Tomislav Futivic
Laurence Gagliani
Edward Gallagher
Andrew Gallo
Chris Gamboni
Xosé Bravo Garcia
Osvaldo Gazzaniga
Kevin Gee
Rodney Gehrke
Radu Cristian Gheorghiu
Greg Giessow
John Gilbert
Serge Van Ginderachter
Greg Goddard
Tiago Goncalves
Ron Goodheart
Octavio Alfageme Gorostiaga
Barry Greene
Jeffrey Greene
Richard Gregor
Martijn Groenleer
Geert Jan de Groot
Ólafur Guðmundsson
Christopher Guemez
Rafael Leon Guerrero
Gulf Coast Shots
Galen Guyer
Sheryll de Guzman
Rex Hale
Jason Hall
James Hamilton
Darow Han
Handy Networks LLC
Stephen Hanna
Martin Hannigan
John Hardin
David Harper
Edward Hauser
David Hauweele
Marilyn Hay
Headcrafts SRLS
Hidde van der Heide
Johan Helsingius
Robert Hinden
Michael Hippert
Damien Holloway
Alain Van Hoof
Edward Hotard
Bill Huber
Hagen Hultzsch
Kauto Huopio
Asbjørn Højmark
Kevin Iddles
Mika Ilvesmaki
Karsten Iwen
Joseph Jackson
David Jaffe
Ashford Jaggernauth
Thomas Jalkanen
Jozef Janitor
Martijn Jansen
John Jarvis
Dennis Jennings
Edward Jennings
Aart Jochem
Nils Johansson
Brian Johnson
Curtis Johnson
Don Johnson
Richard Johnson
Jim Johnston
Jose Enrique Diaz Jolly
Jonatan Jonasson
Daniel Jones
Gary Jones
Jerry Jones
Michael Jones
Amar Joshi
Javier Juan
David Jump
Anders Marius Jørgensen
Merike Kaeo
Andrew Kaiser
Vladislav Kalinovsky
Naoki Kambe
Akbar Kara
Christos Karayiannis
Daniel Karrenberg
David Kekar
Stuart Kendrick
Robert Kent
Robert Kerman
Thomas Kernen
Jithin Kesavan
Jubal Kessler
Shan Ali Khan
Nabeel Khatri
Dae Young Kim
William W. H. Kimandu
John King
Russell Kirk
Gary Klesk
Anthony Klopp
Henry Kluge
Michael Kluk
Paul Knight
Andrew Koch
Ia Kochiashvili
Carsten Koempe
Richard Koene
Alexader Kogan
Matthijs Koot
Antonin Kral
Robert Krejčí
John Kristoff
Terje Krogdahl
Bobby Krupczak
Murray Kucherawy
Warren Kumari
George Kuo
Dirk Kurfuerst
Mathias Körber
Darrell Lack
Andrew Lamb

Richard Lamb
Yan Landriault
Edwin Lang
Sig Lange
Markus Langenmair
Fred Langham
Tracy LaQuey Parker
Christian de Larrinaga
Alex Latzko
Jose Antonio Lazaro Lazaro
Antonio Leding
Rick van Leeuwen
Simon Leinen
Anton van der Leun
Robert Lewis
Christian Liberale
Mark Lieu
Martin Lillepuu
Roger Lindholm
Link Light Networks
Art de Llanos
Mike Lochocki
Chris and Janet Lonvick
Mario Lopez
Sergio Loreti
Richard Lotz
Eric Louie
Adam Loveless
Josh Lowe
Guillermo a Loyola
Hannes Lubich
Dan Lynch†
David MacDuffie
Sanya Madan
Miroslav Madić
Alexis Madriz
Carl Malamud
Jonathan Maldonado
Michael Malik
Tarmo Mamers
Yogesh Mangar
John Mann
Bill Manning†
Diego Mansilla
Harold March
Vincent Marchand
Normando Marcolongo
Gabriel Marroquin
David Martin
Jim Martin
Ruben Tripiana Martin
Timothy Martin
Carles Mateu
Juan Jose Marin Martinez
Ioan Maxim
David Mazel
Miles McCredie
Gavin McCullagh
Brian McCullough
Joe McEachern
Alexander McKenzie
Jay McMaster
Bruce McNamara
Mark Mc Nicholas
Olaf Mehlberg
Carsten Melberg
Kevin Menezes
Bart Jan Menkveld
Sean Mentzer
Eduard Metz
William Mills
David Millsom
Desiree Miloshevic
Joost van der Minnen
Thomas Mino
Rob Minshall
Wijnand Modderman-Lenstra
Mohammad Moghaddas
Charles Monson
Andrea Montefusco
Fernando Montenegro
Roberto Montoya
Joel Moore
Joseph Moran
John More
Maurizio Moroni
Brian Mort
Soenke Mumm
Tariq Mustafa
Stuart Nadin
Michel Nakhla
Mazdak Rajabi Nasab
Krishna Natarajan
Naveen Nathan
Ryan Nelson
Darryl Newman
Mai Nguyen
Thomas Nikolajsen
Paul Nikolich
Travis Northrup
Marijana Novakovic
David Oates
Ovidiu Obersterescu
Jim Oplotnik
Tim O'Brien
Mike O'Connor
Mike O'Dell
John O'Neill
Carl Önne
Packet Consulting Limited
Carlos Astor Araujo Palmeira
Gordon Palmer
Alexis Panagopoulos
Gaurav Panwar
Sujith Madathil Parambath
Chris Parker
Alex Parkinson
Craig Partridge
Manuel Uruena Pascual
Ricardo Patara
Dipesh Patel
Dan Paynter
Leif-Eric Pedersen
Rui Sao Pedro
Juan Pena
Luis Javier Perez
Chris Perkins
Michael Petry
Alexander Peuchert
David Phelan
Harald Pilz
Derrell Piper
Rob Pirnie
Jorge Ivan Pincay Ponce
Marc Vives Piza
Victoria Poncini
Blahoslav Popela
Andrew Potter
Ian Potts
Eduard Llull Pou
Tim Pozar
David Preston
David Raistrick
Priyan R Rajeevan
Balaji Rajendran
Paul Rathbone
William Rawlings
Mujtiba Raza Rizvi
Bill Reid
Petr Rejhon
Robert Remenyi
Rodrigo Ribeiro
Glenn Ricart
Justin Richards
Rafael Riera
Mark Risinger
Fernando Robayo
Michael Roberts
Gregory Robinson
Ron Rockrohr
Graziano G Rodegari
Carlos Rodrigues
Magnus Romedahl
Lex Van Roon
Marshall Rose
Alessandra Rosi
David Ross
William Ross
Boudhayan Roychowdhury
Carlos Rubio
Rainer Rudigier
Timo Ruiter
RustedMusic
Babak Saberi
George Sadowsky
Scott Sandefur
Sachin Sapkal
Arturas Satkovskis
PS Saunders
Richard Savoy
John Sayer
Phil Scarr
Gianpaolo Scassellati
Elizabeth Scheid
Jeroen Van Ingen Schenau
Carsten Scherb
Ernest Schirmer
Benson Schliesser
Philip Schneck
James Schneider
Peter Schoo
Dan Schrenk
Richard Schultz
Timothy Schwab
Roger Schwartz
SeenThere
Scott Seifel
Paul Selkirk
Andre Serralheiro
Yury Shefer
Yaron Sheffer
Doron Shikmoni
Tj Shumway
Jeffrey Sicuranza
Thorsten Sideboard
Greipur Sigurdsson
Fillipe Cajaiba da Silva
Andrew Simmons
Pradeep Singh
Henry Sinnreich
Geoff Sisson
John Sisson
Helge Skrivervik
Terry Slattery
Darren Sleeth
Richard Smit
Bob Smith
Courtney Smith
Eric Smith
Mark Smith
Tim Sneddon
Craig Snell
Job Snijders
Ronald Solano
Asit Som
Ignacio Soto Campos
Evandro Sousa
Fredrik Söderblom
Peter Spekreijse
Thayumanavan Sridhar
Paul Stancik
Ralf Stempfer
Matthew Stenberg
Martin Štěpánek
Adrian Stevens
Clinton Stevens
John Streck
Martin Streule
David Strom
Colin Strutt
Viktor Sudakov
Kathleen Summers
Edward-W. Suor
Vincent Surillo
Terence Charles Sweetser
T2Group
Roman Tarasov
David Theese
Rabbi Rob and Lauren Thomas
Douglas Thompson
Kerry Thompson
Lorin J Thompson
Jerome Tissieres
Fabrizio Tivano
Peter Tomsu Fine Art Photography
Joseph Toste
Rey Tucker
Sandro Tumini
Angelo Turetta
Brian William Turnbow
Michael Turzanski
Adam Tuxbury
Phil Tweedie
Steve Ulrich
Unitek Engineering AG
John Urbanek
Martin Urwaleck
Bart Vanautgaerden
Betsy Vanderpool
Surendran Vangadasalam
Ramnath Vasudha
Jose Luis Couto Vázquez
Randy Veasley
Philip Venables
Buddy Venne
Alejandro Vennera
Luca Ventura
Scott Vermillion
Tom Vest
Peter Villemoes
Vista Global Coaching & Consulting
Dario Vitali
Marc Vives
Rüdiger Volk
Jeffrey Wagner
Don Wahl
Michael L Wahrman
Lakhinder Walia
Laurence Walker
Randy Watts
Andrew Webster
Christoph Wegener
Jd Wegner
Tim Weil
Westmoreland Engineering Inc.
Rick Wesson
Peter Whimp
Russ White
Jurrien Wijlhuizen
William Willaford
Joseph Williams
Derick Winkworth
Pindar Wong
Brian Woods
Makarand Yerawadekar
Phillip Yialeloglou
Janko Zavernik
Bernd Zeimetz
Muhammad Ziad Ziayuddin
Tom Zingale
Matteo Zovi
Jose Zumalave
Romeo Zwart
廖明沂.

Call for Papers

The *Internet Protocol Journal* (IPJ) is a quarterly technical publication containing tutorial articles ("What is...?") as well as implementation/operation articles ("How to..."). The journal provides articles about all aspects of Internet technology. IPJ is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. In addition to feature-length articles, IPJ contains technical updates, book reviews, announcements, opinion columns, and letters to the Editor. Topics include but are not limited to:

- Access and infrastructure technologies such as: Wi-Fi, Gigabit Ethernet, SONET, xDSL, cable, fiber optics, satellite, and mobile wireless.
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance.
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, troubleshooting, and mapping.
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, cloud computing, and quality of service.
- Application and end-user issues such as: E-mail, Web authoring, server technologies and systems, electronic commerce, and application management.
- Legal, policy, regulatory and governance topics such as: copyright, content control, content liability, settlement charges, resource allocation, and trademark disputes in the context of internetworking.

IPJ will pay a stipend of US$1000 for published, feature-length articles. For further information regarding article submissions, please contact Ole J. Jacobsen, Editor and Publisher. Ole can be reached at **ole@protocoljournal.org** or **olejacobsen@me.com**

Follow us on X and Facebook **@protocoljournal** **https://www.facebook.com/newipj**

Supporters and Sponsors

Supporters

Internet Society

CISCO

Diamond Sponsors

Your logo here!

Ruby Sponsors

APNIC FOUNDATION

Sapphire Sponsors

ICANN

Emerald Sponsors

Akamai

amsix

APRICOT Asia Pacific Regional Internet Conference on Operational Technologies

.auDA

DE CIX

EQUINIX

FLEXOPTIX

GÉANT

Google

identity digital

JPRS

RIPE NCC RIPE NETWORK COORDINATION CENTRE

lacnic

LinkedIn

linx

NTT DATA

WIDE PROJECT

VERISIGN

TEAM CYMRU

Corporate Subscriptions

AFRINIC The Internet Numbers Registry for Africa

COMCAST

DNS-OARC

ISC Internet Systems Consortium

IWL

NSRC Network Startup Resource Center

qa cafe

For more information about sponsorship, please contact **sponsor@protocoljournal.org**

The Internet Protocol Journal
Link Fulfillment
7650 Marathon Dr., Suite E
Livermore, CA 94550

CHANGE SERVICE REQUESTED

The Internet Protocol Journal is published quarterly and supported by the Internet Society and other organizations and individuals around the world dedicated to the design, growth, evolution, and operation of the global Internet and private networks built on the Internet Protocol.

Email: ipj@protocoljournal.org
Web: www.protocoljournal.org

Printed in the USA on recycled paper.