The Internet Protocol Journal

May 2026 Volume 29, Number 1

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

I sometimes reflect on the amazing developments in technology that have taken place during my lifetime. In my early childhood, most electronic devices operated using bulky vacuum tubes. The transistor was invented some ten years before I was born, but widespread use of transistors did not occur until the 1960s. Today, chip makers can produce devices with billions of transistors powering devices such as smartphones and powerful computers. When I first encountered the *Advanced Research Projects Agency Network* (ARPANET)—the predecessor to the Internet—in 1976, access to networked devices consisted of terminals operating at speeds ranging from 300 to 1,200 bits per second and the only available and accessible documents were ASCII text files. No images, no audio, and no video. In our first article, Geoff Huston examines how the Internet has evolved from early concepts to worldwide deployment, powered largely by the ever-increasing capabilities of silicon devices. He also discusses the changing role of IPv4 and IPv6 addresses as the overall architecture of the Internet has changed over time.

Domain Names are often linked to branding and intellectual property, and thus they represent commodities that can be bought and sold. In our second article, David Strom outlines the many steps he went through in order to sell his **strom.com** domain.

David Jack Farber served on the Editorial Advisory Board of this journal from its inception in 1998 until our relaunch in 2014. Farber passed away on February 7, 2026, at the age of 91. We are grateful for his many contributions to the success of this journal and to the Internet in general. We will miss him.

All 96 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 3,600 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

Internet Evolution

by Geoff Huston, APNIC

I work as the Chief Scientist at the *Asia Pacific Network Information Centre* (APNIC), the *Regional Address Registry* (RIR) for the Asia Pacific, which provides services related to the management of IP numbers (Addresses and *Autonomous System Numbers*) to networks in this region. I sometimes like to lift my head above the address parapet; elevate the typical RIR policy conversations above the day-to-day mundanities of address allocation policies with their vocabulary of address block sizes and needs-based justifications, fairness, and efficiency; look more broadly at the context of the industry we operate in; try to gain an understanding of where we are right now; and speculate on where it's all going.

The Rise of the Internet

The Internet is the outcome of one of the most prodigious marriages in all human history, namely the marriage of computing and communications. Since the inventions of the transistor in December 1947 and the integrated circuit in 1958, the technology world—and the broader realm of human societies—has been fundamentally changed. Until this time all human endeavours had been limited by their geography. The industrial age and the explosive introduction of the railway in the mid-19th century heralded a significant societal change as nation states shifted their foundations of wealth and power from land and agricultural activities to industrial production, and the scale of activities lifted from cottage industries to enterprises that operated at a regional or national scale. Facilitating these changes was firstly the telegraph and then the telephone, allowing these new industrial enterprises to project their power and influence across larger expanses, harnessing greater volumes of production and increasing their wealth. When we added computers into the communications realm, the rate of change increased dramatically. It took us a decade from the invention of the transistor to that of the integrated circuit, and a further decade to reach the point where computers passed through the transition from esoteric research endeavours to an essential tool for data processing and communications, and by then we were on a path that in retrospect was unstoppable!

In the late 1960s, Bell Labs programmers Ken Thompson and Dennis Ritchie devised the *Unix* operating system. It was written in the *C* language, a so-called high-level language, and its associated compiler could produce assembler code for a variety of computers. It was one of the first of the "open" operating systems of the age. This openness was not exactly an unforced choice: Under a 1956 consent decree in settlement of an antitrust case, Bell System (the parent organization of Bell Labs) was forbidden from entering any business other than common carrier communications services, and was required to license any patents it had upon request. Unix could not, therefore, be turned into a product. Bell Labs shipped the source code Unix system for the cost of media and shipping to those who asked, allowing universities and other organizations to modify and extend it.

This open model, both of the code itself and the refinement and development of the system, facilitated the development of many variants in the ensuing years, including the highly influential *Berkeley Software Distribution* (BSD). Open software has formed the backbone of the computer industry ever since.

In 1973 Bob Metcalf, working at the *Xerox Palo Alto Research Center* (PARC), published a memo describing *X-Wire*, a 3 million bits-per-second (Mbps) common bus local-area network, which became known as *Ethernet*[0]. It was notable because Ethernet was the absolute minimal form of computer networking. It really was just a wire. It had no internal switch, no packet framing, no clock, no controller, no network state. Nothing! Just a wire. And in many ways Ethernet changed the way we thought about computer networks and communications architectures for computers. There was no technology inside the wire. It was the simplest network you could ever have. It was a wire. Everything happened in the connected computers at the network edge. And that's why Ethernet took off, because every conventional function of network control was pushed off the network into the connected computers and implemented as a collection of distributed algorithms. The network itself, this wire, was just transmission—and only transmission. "Dumb network, Smart devices," as we constantly reminded ourselves. What this also meant is that the money to operate the network wasn't the network's problem. That's up to the computers that connected to this network. The connected computers were operating the software, so the network was just a packet transmission medium (Figure 1).

Figure 1: X-Wire Framework

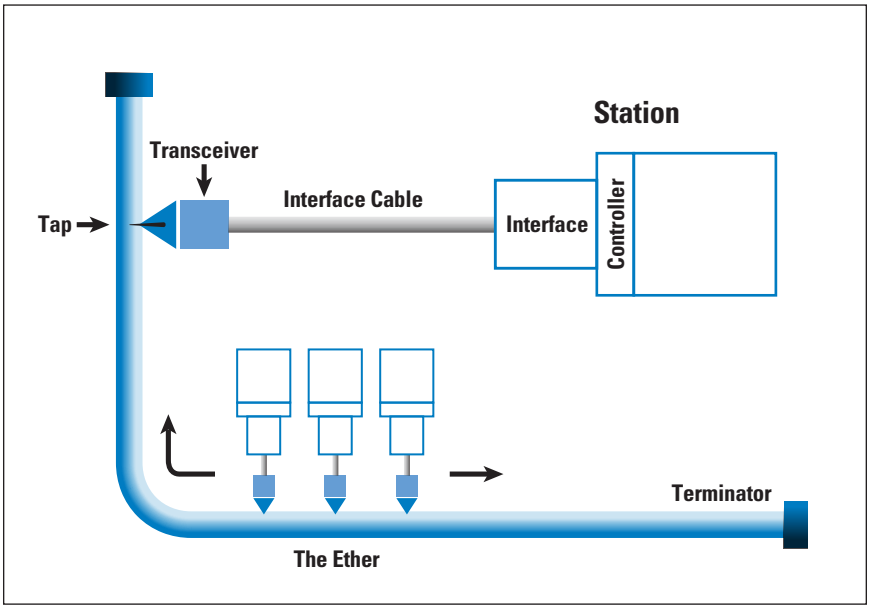

So, the collective "intelligence" for an Ethernet *Local-Area Network* (LAN) was placed into the computers that connected to this common wire. All Ethernet packets were self-clocked by using a 64-bit preamble to allow all receivers to synchronise their receiving data clocks against the data rate of the packet being transmitted.

Packets were between 64 and 1,518 bytes in total length. There was no centralized contention control—each station used a *Carrier Sense Multiple Access with Collision Detection* (CSMA/CD) protocol to negotiate exclusive sending access to the wire, and all connected computers received all packets that were passed onto the wire. It used a unique station addressing plan in the form of a 48-bit *Media Access Control* (MAC) address that is still with us today! But most importantly, like Unix, it was an open standard specification.

Figure 2 takes us back to 1977, with the release of the Vax 11/780 computer, made by the *Digital Equipment Corporation* (DEC). It was a medium-size mainframe computer capable of executing 1 *Million Instructions per Second* (MIPs). A very popular computer, it was used in many corporate environments as well as universities and research institutions. It signified a visible change in the computing industry where useful computers were physically shrinking, while at the same time their processing capabilities and storage capacity were increasing dramatically.

Figure 2: DEC Vax 11/780.

If you look at other 50-year-old technology artefacts, many of these items don't look all that different from what we use today. Yes, cars today are lighter, more efficient, and come with a whole lot less chrome, but they still have a wheel at each corner and some internal engine for propulsion. And I guess that 50-year-old cars are still around today, and doubtless some of them are still on the road. But this is not the case for computers. There are very few 50-year-old computers left in today's world, and those that still exist are in museums. Why has the computer industry been so ready to cast off its past lives and embrace change so readily?

The essential difference here is that the pace of *Moore's Law* in the improvements in the production of integrated circuits on silicon wafers has been truly miraculous. Originally an observation by Intel co-founder Gordon Moore in 1965, Moore's Law says that the number of gates on an integrated circuit doubles approximately every two years, while the cost of fabrication of the chip has increased by a far lower factor, if at all. The cost and quantity of compute capability have improved inexorably over the same 50 years (Figure 3).

Figure 3: Moore's Law[1].

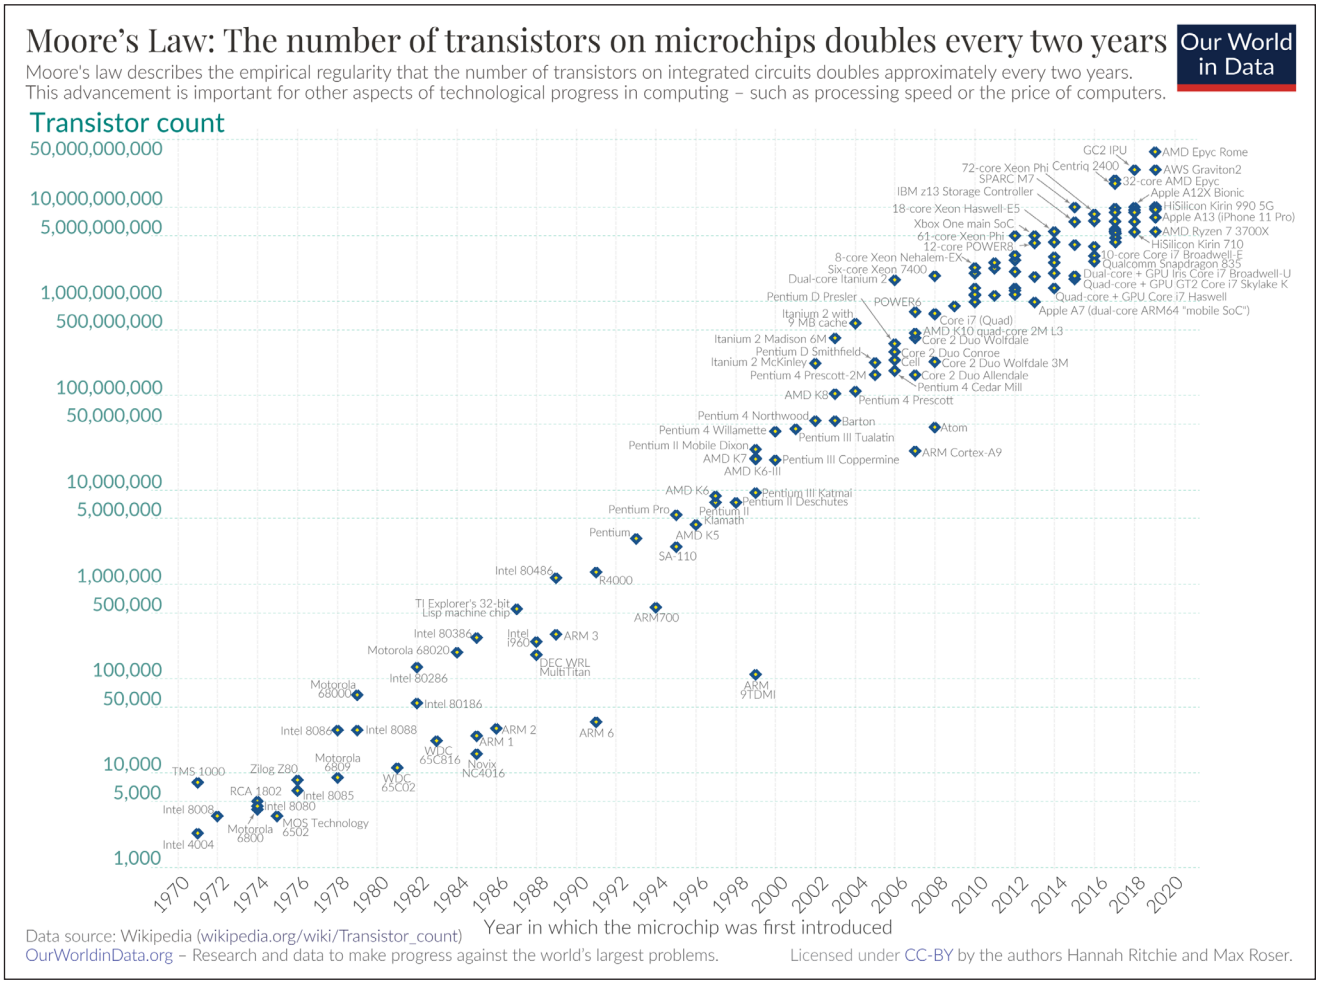

It is often hard to come to terms with exponential growth. It has its own momentum and raft of unintended consequences. Initially, this growth was all about making these lumbering mainframe computers faster, smaller, and cheaper. The mainframe computers of the early 1960s that took up an entire data room with multiple racks of circuitry, power, and air conditioning changed every few years. These machines had assured obsolescence from the moment they were made, and their useful service life was best measured in units of months rather than decades!

This dynamic pace of evolution of computing capability has impacted many aspects of the computing environment. Network protocol design was certainly part of the impact zone of these improvements in processing capability. One of the objectives of a computer network was to allow any connected computer to communicate with any other. In order to achieve that goal, each computer needed to have a unique identification field, or an "address" to distinguish it from all the other connected computers. The proprietary network protocol used by the VAX Systems was *DECnet*, and at that time Phase 3 of this network architecture was in widespread use. It used a 16-bit address field, allowing a maximum of 65,535 computers in a single DECnet environment. In the late 1970s the concept of a single network populated with tens of thousands of these large lumbering computers was an unaffordable fantasy. But that thinking did not factor in exponential growth through Moore's Law, and the associated exponential drop in size and cost. The computing environment rapidly shifted in the 1980s into personal computers, and the multi-million-dollar mainframes of the 1970s with their specialised environments of highly curated clean rooms transformed into small boxes that found a home on every desk, in both the workplace and the home. What was an esoteric fantasy of running out of 16-bit DECnet addresses in the 1970s was an obvious reality within just ten years.

But not every computer networking protocol had such a severely constrained design. The folks working on the *Internet Protocol* (IP) took a far more radical step and used a 32-bit device address architecture. Instead of an inbuilt ceiling of some 65,000 connected devices, the IP protocol had a ceiling of some 4.3 billion connected devices! This thinking was crazy, even during the period when computers escaped from their clean rooms and invaded offices and homes. This number was larger than the world's population at the time. This computer protocol matched the inexorable progress of the expansion of the computing resolution.

We now had a triumvirate in the combination of the advances in the speed, size, and cost of computing, an openly available, highly efficient operating system for these computers in the form of Unix and a freely available computer networking protocol in IP that had this massive inbuilt capacity through its crazy decision to use a 32-bit addressing architecture!

And that was why the IP design was so revolutionary. A few years after that photo in Figure 2 was taken, a few hundred thousand of these Vax computers were in the field. And we looked at the IP address space—with some 4 billion addresses—and said to ourselves: "What's your problem? We're here forever!" Awesomely wrong, as it turned out!

Moore's Law continued to play havoc with this industry. At the same time as the VAX 11/780 was released, over on the West Coast of the U.S. the Apple Corporation released its first Apple computer. This device was never going to challenge the mainstream computer industry.

It cost hundreds of dollars, not millions, it lived on the desktop, and it played games. But Steve Jobs, Apple's founder, appreciated the true effects of Moore's Law. In the search for untapped markets that could absorb the production capacity of the silicon chip industry, Apple was a pioneer in the consumer market and applied design principles to personal computers. This computer was not only at home, but it also looked like it was truly *part* of that home.

In addition, Steve understood more than just the effects of Moore's Law in terms of cheaper and smaller; he understood Moore's Law in terms of using that computing capacity to elevate his machines to interact with its human user on human terms. Apple developed the *Macintosh*, a computer with an intuitive look and feel, to the extent that it did not need a geeky user manual. (Coincidentally, this graphical interface was also developed at the Xerox PARC, and while Xerox was highly influential in these landmark inventions, the company stubbornly stuck to being a photocopier company!)

When you powered up an Apple Mac, it displayed a friendly "Hello!" The human on the other side of the keyboard didn't need to memorise some arcane Job Control Language to make this machine sing and dance. It used a part of its compute capability to define an interaction with its users that was phrased entirely in human terms. And in the consumer market it worked! The personal computer market quickly outpaced the mainframe market and dominated the entire computing environment.

Something else was also happening then. In these early days of computer networks, we viewed the network in the same terms as we viewed the telephone network. Human users of telephones both spoke and listened. The telephone network was symmetric and invisible. In computer networks the connected devices both provided digital services and accessed the services of others. Connected computers were both client and server in the networking environment. But that's not really thinking about Moore's Law in a useful way. Sure, there was a driving need for more and more personal computers, but these computers were not both client and server computers—they were positioned as simple clients. Segmentation of connected computers emerged into dedicated client and server roles. The large mainframe computers didn't disappear; they were used to provide services to the millions, then tens of millions of users behind these client personal computers.

It was no longer a computer network built in the image of a telephone network, but a network that was modelled on broadcast television, with slightly better programming control that allowed every client to perform its own programming. Users at home were not providers, and they did not actually want to host services, and certainly did want these service delivery platforms to be located in their homes. They just wanted the computing equivalent of the television set. Consequently, we constructed the emerging Internet network of the late 1990s around that model of clients and servers.

At the time this plan was convenient because we were still building the Internet on the underlay of the telephone network. We applied this entire asymmetric behaviour of clients and servers back into the network architecture of the Internet itself. If you were a client, you couldn't do much at all, but you could make contact with servers. The connected edge client computer was the thing that simply said, "Let's look at the data out there." The dial-up world of the 1990s and the *Digital Subscriber Line* (DSL)/Cable Modem world of the 2000s in the last-mile access network were a good fit to the demands of client/server networking at the time, and the Internet rapidly expanded by repurposing the existing last-mile infrastructure, avoiding the need for hefty initial capital investment by *Internet Service Providers* (ISPs).

Millions of people were eager to own their personal computers and use the Internet to access digital services. The industry continued to grow. By around the year 2000 we saw the use of specialised data centres that were intended to coalesce all these servers into a common curated environment with power, cooling, and a couple of human attendants to scurry about and replace dead or dying server units. This server world also saw role specialisation. There were web services, mail services, data services, archival services, and so on.

Compared to the current AI-scale data centre designs, this activity was small-scale. These Internet service data centres were only a room or two with a few hundred megawatts of power. All of the online services were now crammed into these dedicated service delivery points. As we continued to expand into the consumer space, we turned to large-scale "network exchanges" and "network peering points"[3a,3b] to service the concentration points where data intersected with service delivery capability. Clients never exchanged data with other clients. Even transaction was a transaction between a client and one or more servers.

And then the network moved into the whole thing about network traffic engineering. This transition was interesting in that the silicon industry was working like crazy hard to expand the client base of the network, but at the same time we were spending a lot of time and measure to control and even ration the use of the network to ensure that the limited resource, network capacity, was used fairly and efficiently.

This network was now a long way away from being "just a wire." It was replete with hardware and software to perform all kinds of functions relating to resource control, variable service response, defensive measures, and of course, data collection. The progress of Moore's Law enabled addition of further functionality as the processing capability of these service delivery platforms increased. But this progress of silicon chip capability extended to more than the capabilities of the data centre services and in-network middleware. When you shine a laser light through a fibre optic cable and modulate the light to carry a digital signal with on/off keying, then you can work with a simple light detector.

If you want additional data capacity from a fibre circuit, then the logical approach is to use the same signal modulation techniques we previously used for analogue modems, namely using phase and amplitude modulation of the light signal. While this capability can radically improve the carrying capacity of a fibre system, the practical limits of this approach are strongly dependent on the signal processing capabilities used at either end. The reason why you can increase to almost a terabit-per-second of signal through fiber optic cable is not solely because you are doing anything radically different with glass fibre, but also the increased capability of the *Digital Signal Processors* (DSPs) you're using at either end. If you are using three-nanometre chip technology, it appears that you can achieve some 800 Gbps from a single polarisation of a light signal through fibre. It's not necessarily easy, but it's achievable. At the same time, we've lifted the common network bearer capacity to 100–200G with these latest generations of DSPs.

A Network of Abundance

These developments have heralded a huge shift in the role of the network. When transmission capacity was a scarce resource, rationing access to resource across competing demands was the role of the network. Rationing scarcity is a high-value task, as in economic terms the controller of the resource can demand a "scarcity premium" from competing users. But if capacity is abundant to the extent that it overwhelms all demands, then the network operator's role shifts into a basic undistinguished commodity brokering role. We've seen this shift in the commercial world, where the share value of network carriage operators has plummeted in line with this shift from scarcity to abundance. The response from the network operations sector has been an effort to add further functionality to the network, in the form of variable service responses, network segmentation, and customisable forwarding functions, but to date this effort has increased costs in the network while at the same time failing to attract sufficient interest from the network client base. The same basic improvement in technology capacity also means that connected devices are increasingly capable of managing their own requirements for service responses from an amply dimensioned abundant network. In this environment of network abundance, the network service itself quickly becomes an undistinguished commodity good and is priced accordingly.

The silicon industry then turned its attention to handheld and embedded devices, and once more, it overachieved! Billions of devices were added into the Internet in the short span of a couple of years as a result. The cumulative result is that we are now in a digital world that is defined largely by abundance and scale: abundant processing, abundant storage, and abundant carriage capacity.

What happens in an environment of abundance? This question is very relevant, as abundance drives today's Internet. Scarcity has driven this industry until today—and now we're in a different industry where almost everything is being commoditized and is becoming overwhelmed with capacity. And where we used to use pricing to ration a limited resource when performing the distribution function, we are now looking for a new model. Pricing isn't an answer to this new challenge, so if you're a carrier life is now grim. The carrier's historical role is to take a single resource and share it amongst the folks who want to use it. But that paradigm is over, and the underlying value of the network infrastructure is extremely low. Didn't T-Mobile sell the Sprint wireline network to Cogent for a single dollar? They probably paid too much.

Abundance appears to have destroyed the network completely! The role of the network was a "just-in-time" service that connected users to service on demand. This model has been largely replaced by *Content Delivery Networks* (CDNs), whose delivery model is to position replicas of the service content across a large set of service points. This model could be characterised as a "just-in-case" model of service delivery, where content is pushed out to the edges, close to potential consumers, just in case they ask for it.

This model has had its own ripple effect. As we move content and service closer to users, the network becomes smaller. We are eliminating distance as the dominating characteristic of the network. We don't grab content from far away any more. The packet-miles, or the distance a packet travels from the server to the client, are now tiny. When you shrink the network from the circumference of the globe or across a continent down to the dimensions of traversing a single city, then everything is easier, faster, and cheaper. If you can build an internet that looks and performs like a Local-Area Network, then where is the scarcity? Where are the performance issues? They are largely gone. When you shrink distance in networking, everything else becomes easier.

Bigger, Faster, and Cheaper

Today's distance-shrunken Internet is, oddly enough, bigger. We are talking terabits of capacity in the core systems and gigabits out to the edges, in both last-mile fibre and last-mile radio. We just publish all the content in all the places, all at once, because that's what abundance says. All of this content deployment shift has allowed us to move to a network that is phenomenally big because it's not "distance-big," it's "capacity-big" because when you shrink distance you can improve protocol performance and decrease cost.

This network is now fast. Really fast. Not only do I get a gigabit per second to my house, but I can actually pull down content with a cooperating server or two at that same line speed. What enables this increased performance is shorter network distance.

Network transport protocols are feedback-controlled protocols, and when you bring the two ends closer together you increase the number of potential feedback cycles per unit of time, and this increased information flow provides better guidance to the transport protocol to optimise efficiency of delivery.

And yes, it's cheaper. If I took a gigabit-per-second Internet service and pushed it back to the currency unit of 1975, fifty years ago, then the $100 per month that users pay on average for such a residential service would cost $17 of the 1975 dollars. And not only is this service cheap by historic standards, but its reach has also extended all the way to the most remote corners of the surface of the earth, with the advent of low-earth-orbit satellite-based services.[4]

We seem to have achieved the elusive triple of better, faster, and cheaper. The engine that has got us here is this continuous refinement in the fabrication of silicon chips.

The End of Moore's Law?

So far, we've looked at the upbeat story of the past fifty years. But while Moore's Law has carried the burden of the progressive improvement in this industry for the past sixty years, are we reaching the end of this journey? Down on the silicon chip we're down at feature sizes of 3 nanometers, and it looks like 2 nanometers is achievable. How much further can we take this continuous improvement in silicon chip manufacture? Or, to put it more bluntly, when will we reach the end of Moore's Law? To put that into context, a silicon atom has a diameter of 0.23nm. The implication is that a chip feature of 2nm in size is only 10 silicon atoms wide.

Can we go to 1-nanometer feature sizes on silicon chips? Well, we can try. We are using extreme ultraviolet rotation already to produce a wavelength of 13.5nm for chip lithography, and the process is extremely complex. Even if we can achieve a reliable and scalable lithography process, there is also the problem of electrons, and let's assume they're particles even though there is the duality of quantum mechanics. In small enough spaces electrons have a mind of their own. It is hard to get this chip feature size even smaller using today's planar integrated circuit technology. That thing we did in 1957 to create an integrated circuit has been progressively refined year after year, and we've made the features smaller and smaller as a result, but it's possible that this process of continuous improvement will grind to a halt. Nobody is announcing a 4 trillion gate silicon chip this year, and I'm not sure they ever will. This technology as it stands is close to that physical edge of impossibility, and it appears that the next innovation needs to be really off-the-wall clever and do something that none of us have thought possible.

That's not the only limitation we've encountered in chip capability. The clock speed of processors topped out in the early 2000s. Processors may have more gates, and CPUs may have more cores, but the clock speed of the processor has remained pretty constant since then (Figure 4).

Figure 4: CPU Clock Rate[2].

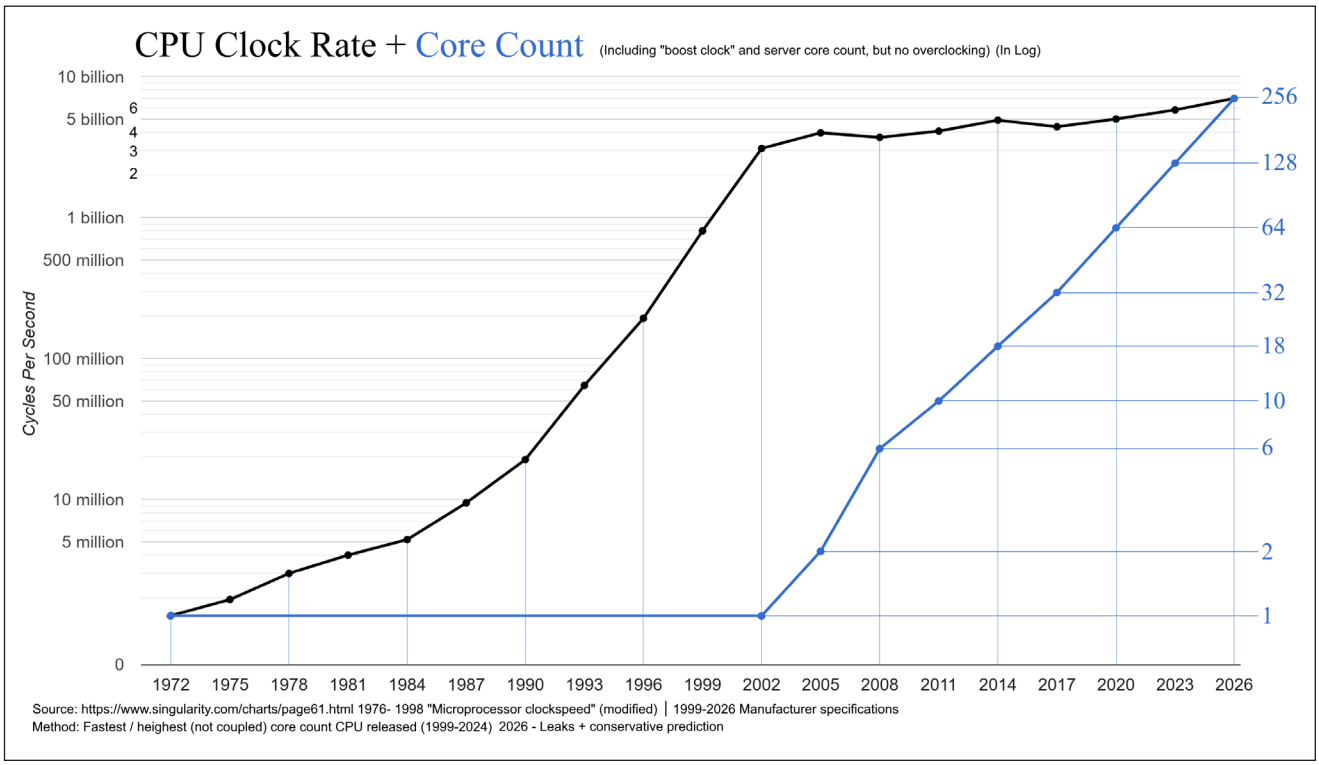

It is way too early to proclaim that Moore's Law is over. But perhaps it would be wise to contemplate what it means for us if Moore's Law becomes the gift that has effectively stopped giving.

IPv6 and IP Addresses

How does all this information relate to IPv6 deployment?

The effort of the early 1990s intended to provide a path through the situation where the population of connected devices exceeds the capacity of the 32-bit address plant. The response, a redefinition of the IP protocol using a 128-bit address field, preserved all the behaviours of the Internet, but the lack of any form of backward compatibility implied that all network users and operators were exposed to the marginal cost of adopting a new network protocol, but in fact gained none of the marginal benefits that a completely different network protocol might offer. It was still just IP.

The alternative option, which was also available, was to deploy *Network Address Translators* (NATs)[5] in the network infrastructure. To make up this shortfall in IPv4 addressing, NATs were adding greater processing capability into the network infrastructure.

A client-side response, NATs are viable in the long term only if we can scale processing efficiency in line with demand growth. The purpose of NATs was bridging that gap between the capabilities of the technology we had and the size of the network that we wanted to deploy. As the network grew, the size of the gap increased, and the cost in terms of infrastructure complexity just got higher and higher. (Figure 5).

Figure 5: Bridging the IPv4 Address Gap.

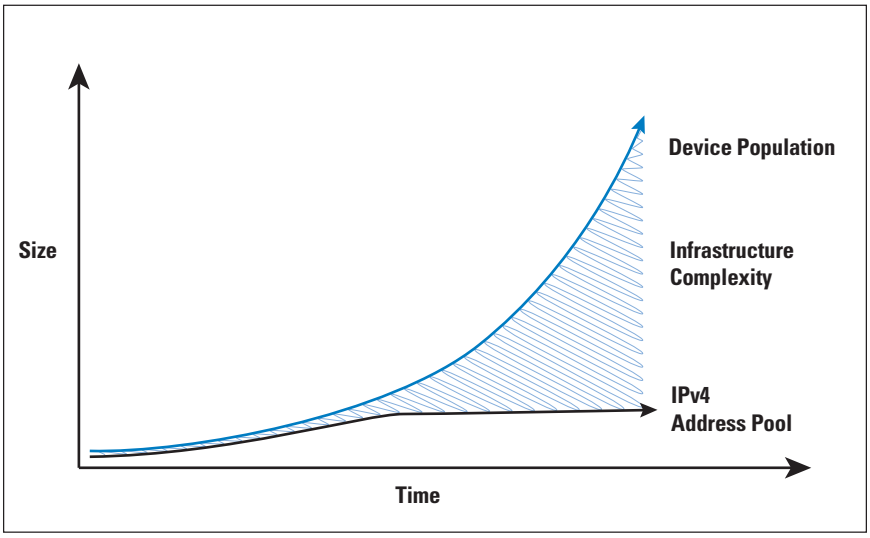

At some point this embedded processing in the network becomes problematic. It's a cost element, and costs in a commodity network are intolerable in the long run.

The NAT response to the IPv4 address exhaustion situation gives us little leeway to reduce network complexity. That means if you want to continue to deploy digital services at scale, contain cost escalation to keep the service affordable, and improve network robustness and security, then you have few choices left other than to reduce the network complexity burden. Deploying IPv6 is one obvious response as a way of achieving this goal.

Who received this memo? India got the memo in 2017 and rolled out a V6 network across a billion users in 12 months. That's getting the message at the scale and speed that they needed to achieve to be a part of today's world. China is facing a similar set of problems, but it is proceeding at a more leisurely pace with their program of IPv6 adoption (Figure 6).

There's another tension going on here: the vertical tension inside the protocol stack. All the value is moving up the protocol stack into the application layer. The application is now the service.

Figure 6: Daily Plot of IPv6 Deployment in China and India.

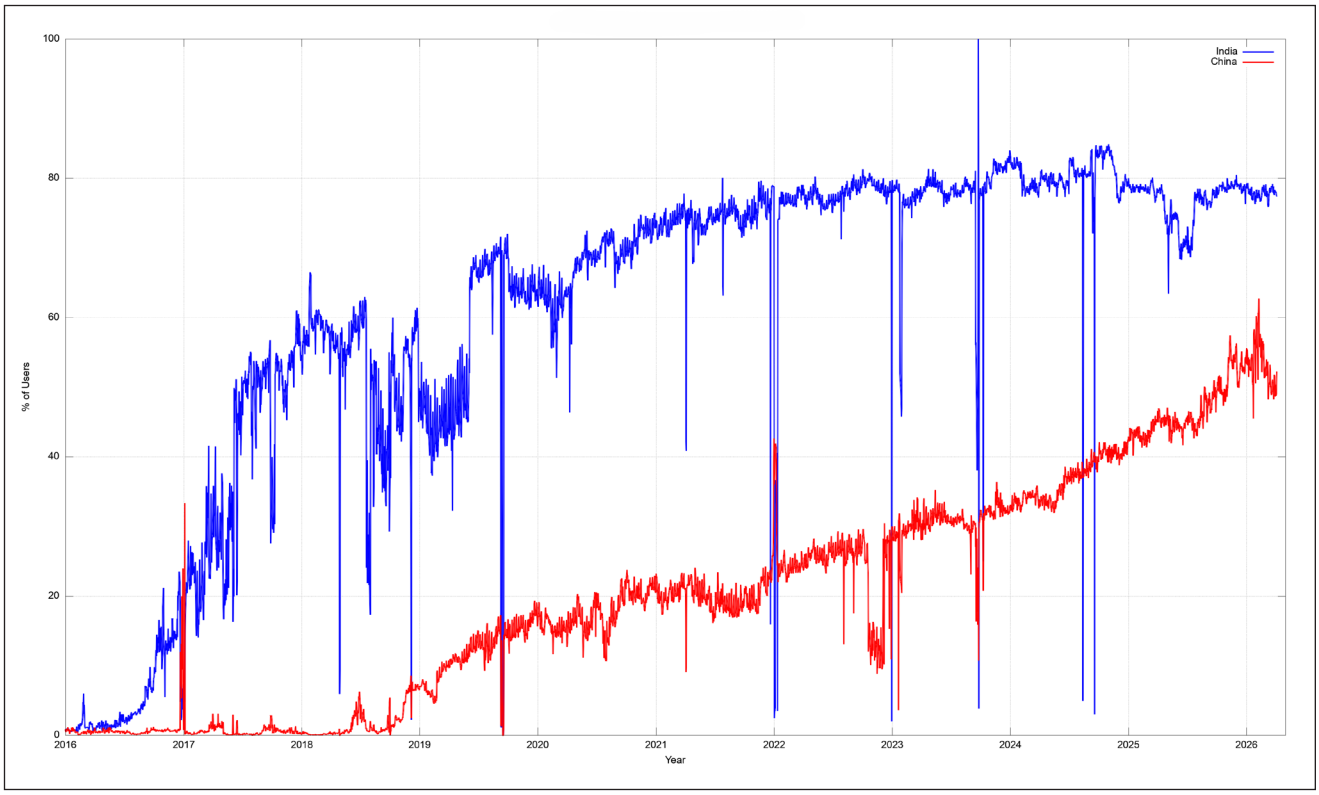

The client/server network needs accessible rendezvous points for servers, or to be more correct, for services, but not for clients. In the public Internet clients do not need to be associated with a persistent IP address, and the QUIC[6,7] transport protocol takes it a step further with the observation that the network can change the IP address of the client at will, and as long as the address assignation is stable for a couple of round-trip time intervals, all will work! For server IP addresses persistence is not a strict prerequisite either. As long as it remains usable across some *Domain Name System* (DNS), resolver cache lifetime IP addresses for servers need not be stable for very long. The reason here is because it's the DNS namespace we use as the persistent service identity schema for the Internet. Our mechanisms for service authenticity and channel security are based on persistent service names, not on their IP addresses.

Therefore, I can't help but ask: Do IP addresses matter? I have to observe that, yes, IP addresses do matter. But I must qualify that to add that they matter at the moment, but they are not necessarily invariant properties going forward. IP addresses are not persistent identity tokens, nor are they useful location tokens. They are useful tokens to pass into the routing system to assist packets to be passed to their intended destination, but even this is not quite what it seems.

For example, the *Multiprotocol Label Switching* (MPLS) approach places a wrapper around an incoming packet that identifies to the network the desired network egress point and a path to reach that point.

On its transit across that MPLS network, the inner IP protocol destination address is not used to forward the packet. So more accurately, individual networks use addresses to determine where to evict each packet from the network!

An extreme view is that IP addresses are ephemeral tokens to map from the service layer of named services to the underlaying layer of packet forwarding, and that's about it. Yes, we've attached a whole set of ancillary roles about reputation, accountability, location, and reporting to an IP address, but perhaps these additional uses of addresses were making assumptions about the role and persistence of addresses that are unwarranted in the longer term.

We are not done. The silicon chip industry is still operating at full throttle, and there is the expectation that all of these processors will end up in Internet-connected devices in one form or another. We need to scale the network further. We have to scale the network further. It's service-level scaling that's the challenge, right up at that top level of trying to orchestrate the behaviour of individual network components to generate coherent outcomes.

Do we know how we will achieve this solution? Not really. But don't forget that this global enterprise is not centrally planned. Noone is in charge. No single country is in charge, despite some persistent suspicions about the role of the United States in the background. But such suspicions are more in the realm of conspiracy theory fodder than a clearly established national role. If there is no single country, then is it a group of countries? There is no specific international treaty to provide oversight for the Internet, no international convention or international protocol. There is no imposition of control over the Internet, or even a clear exposition of direction. The Internet is a diverse collection of markets, large and small. Markets are a means of orchestrating the behaviour of actors in producing coherent outcomes in a manner that also attempts, often imperfectly, to maximise efficiencies along the way. However, markets do not necessarily behave rationally—or even predictably. When we gaze into the crystal ball of the Internet and attempt to make such prognostications about its future, the vision is made as if through a glass darkly.

References and Further Reading

[0] Mikael Holmberg, "The History and Future of Ethernet," *The Internet Protocol Journal*, Volume 27, No. 1, March 2024.

[1] Wikimedia on Moore's Law:
https://commons.wikimedia.org/w/index.php?curid=98219918

[2] Wikipedia article on Clock Rate:
https://en.wikipedia.org/wiki/Clock_rate

[3a] Geoff Huston, "Interconnection, Peering and Settlements—Part I," *The Internet Protocol Journal*, Volume 2, No. 1, March 1999.

[3b] Geoff Huston, "Interconnection, Peering and Settlements—Part II," *The Internet Protocol Journal*, Volume 2, No. 2, June 1999.

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

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

[6] Geoff Huston, "A Quick Look at QUIC," *The Internet Protocol Journal*, Volume 22, No. 1, March 2019.

[7] Geoff Huston, "Comparing TCP and QUIC," *The Internet Protocol Journal*, Volume 25, No. 3, December 2022.

[8] Vint Cerf, "A Decade of Internet Evolution," *The Internet Protocol Journal*, Volume 11, No 2, June 2008.

[9] Geoff Huston, "A Decade in the Life of the Internet," *The Internet Protocol Journal*, Volume 11, No. 2, June 2008.

[10] Geoff Huston, "Twenty Five Years Later" *The Internet Protocol Journal*, Volume 26, No. 1, June 2023.

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

[12] Tejas Suthar, "IPv6 and MPLS," *The Internet Protocol Journal*, Volume 8, No. 2, June 2005.

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

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

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**

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/

Selling My Domain After 30+ Years

by David Strom

Back in 1993, I took the radical (at the time) step of requesting a new domain, **strom.com**. I say request because back then there wasn't any actual "purchase"—the Internet was still relatively unknown by the general public, and all it took to become master of your domain was a simple email request, which was satisfied within minutes. Let us pause to remember and honor those simpler times.

Over the years I have used that domain for my own promotion and my own business, which is mostly freelance writing magazine articles, such as those for *The Internet Protocol Journal* (IPJ), and giving speeches. But the time has come for me to sell the domain. And for those of you who are thinking about selling your domain, here is a guide. If you are using the domain(s), this process isn't simple or straightforward. If you have virgin domains untouched by any content, the process is less complex, but you will still need some guidance. No matter what your situation is, my goal with this article is to help you think things through before you finalize the sale of a domain.

You may, after reading what follows, think it is all too much trouble. That's fine: knowledge is power. Or you may think that there is some psychological or financial or some other benefit. That's fine too. You may realize that you are more organized than I and will have an easier time getting your digital life together for an eventual sale. Whatever your reasons, here we go.

This isn't the first time I've sold my digital assets. In 2020 for IPJ, I wrote about selling an unused IPv4 address block.[1] The process for selling a domain is both more complex and also depends on your particular situation: It depends on what you are doing with your domain, and whether there is sufficient market value to motivate the sale. I was lucky enough to have a short last name, to use it only for my domain, and to have chosen a dot com, which is still the most marketable suffix. For years I regretted that I didn't own the domain **davidstrom.com**—that would have been a better choice for my brand, but not as fungible as a domain with just my last name. Why? It turns out Strom is used in a variety of non-English languages and a variety of businesses.

Finally, a note of caution. I use a lot of *Google Workspace* services, so this guide is geared in that direction. Your mileage may vary, depending on the mix of service providers and other digital things you use. I will include a few comments about Microsoft 365 account migration where appropriate. But wherever you have your digital footprint, the decision and workflow process should be mostly similar (and hopefully more streamlined than what I went through).

The process has several steps, some of which can happen concurrently during execution:

- Figure out who is going to list and eventually sell your domain, whether that is yourself, your ISP, or a third party.
- Inventory your various digital assets that are part of your domain (web, email, servers, and online services) or that depend on it. This step took me the longest (several months!), and while writing this article I remembered several corners of the Internet that I needed to inspect—or revisit—to ensure that I could still access them after **`strom.com`** went away.
- Select and then set up your new domain.
- Begin (and hopefully finish) the data migration process.

Figure Out Who Is Going to List and Eventually Sell Your Domain

Before you get any further along in this process, you need to make a realistic appraisal of your digital assets. Do you expect to make some (or a lot of) money from your domain sale? That seems obvious, but perhaps it is for other reasons (your brand has changed, you are parting ways with your company, starting your own business, etc.). So that begs the next question:

How can you determine what the expected value of your domain name might be? The short answer is you can't. Many of the Internet providers offer calculators. That word is somewhat deceptive, because it implies an accounting of sorts. But if you look at three of them[2], you will find a wide variation in expected value. When I last checked what **`strom.com`** was worth, I got answers that ranged from $24,000 to over $300,000. That may give you false hopes, or it may give you some encouragement to sell your domain. My domain was sold in the middle of this range. The issue with these calculators is that they won't agree because (unlike tracking your finances, say) there is no easy nor consistent way to do the math.

Should you use a brokerage service? Just like for IP address blocks, there are brokers for domain purchases. A good place to start learning about brokerages is on **`VPN.com`**[3]. I suggest you think about a domain broker in two contexts: whether they will find you a buyer, and whether they will help you structure your sale. You might think that a domain broker is like an IP address broker, or even a real estate broker, but they are very different. Some of them are good at searching for potential buyers because they use your name in their businesses, or have some relationship with your name.

I found it is a difficult world to evaluate and figure out, because buyers and sellers have different metrics for how they choose domain names. Nevertheless, people buy and sell domains every minute, and have lots of reasons for choosing a new domain name. But that doesn't mean there is a marketplace similar to selling a house or an IP address range.

Both of the brokers I contacted suggested that the buyer would pay their commission fee, but that is by no means standard. Indeed, I got the feeling that the domain brokerages weren't as well structured as the IP address brokers, perhaps because an IP address seems like more of a commodity. There is no (at least to my thinking) intrinsic value of a low-numbered IP address, for example.

I interviewed two brokers who had widely different business practices. One charged a fee twice as much as the other. One was willing to provide references and one wasn't. But eventually I ended up selling it to a private party that contacted me directly. Both of the brokers I initially looked at were widely optimistic on what they could get for my domain, mirroring the range of online calculators.

The final step in this process is using an escrow service to collect your money and provide final control over your domain's assets. I will have more to say about all this later.

Numerous businesses provide this kind of escrow services, but I found **`escrow.com`** to be first-rate, and I used them to finalize my sale. I also engaged a lawyer to vet the agreement that I signed with my buyer, something you probably want to do even if you have a broker representing you. You can specify who pays the escrow fee (which they are very transparent about) when you set up the sale on their platform.

Inventory Your Assets

So assuming you want to proceed, let's start with understanding what assets will be attached to your existing domain. As I said earlier, if you have been paying to register an empty domain for the past some years, then you can skip through this section, but for the rest of us, you will need to do an inventory of your digital assets. This sounds simple, especially if you are a small business, but we don't really think about our digital footprint until we start examining its reach and depth. In my case, I have two web-based main content collections: a regular website that I have maintained for decades of aging and static HTML pages, and a separate server running a *WordPress* blog that I keep up-to-date. Both of them are maintained by one Internet provider (**`Pair.com`**). This registrar is separate from my domain's registrar (*GoDaddy*) and from my content delivery provider (*Cloudflare*). I have kept them separate for decades, and that introduces a bit of complexity.

But the reach is where things get complicated. I run a mailing list server using another provider (which is mostly independent and outside of **`strom.com`**). I use various Google Workplace services (*Voice*, *YouTube*, *Calendar*, *Groups*, and a few others), and each requires a careful assessment of whether or not I want to migrate them to the new domain or make backup copies for my archives. In my case, because I was a long-time Google user, I didn't pay for these services, and I wasn't willing to migrate from a free service to a paid subscription.

You may not be as penurious as I, or you may want to pay more to stay out of the Googleplex entirely. At last count, Google Workplace offered more than 60 different services, and as you'll see in the migration section, many have oddball workflows to move your content from old to new accounts.

I began this inventory process about nine months before the actual sale of **`strom.com`**. But I wasn't really going about this inventory in any organized way, so keep that in mind as you read through the rest of my process. And it wasn't a full-time job, either, but I fit in tasks between producing actual work product for my clients. Eventually, I created a spreadsheet that listed all the various domains and how they were managed, more than a dozen different WordPress sites that I have created (and abandoned) over the decades, and email addresses that I used across the vast Strom empire. You might want to start such a document of your own, or use something better.

Set Up Your New Domain

Before you begin the migration of your data, you should pick your new domain and get the place set up to receive your existing digital content. If you have a second domain that already exists (which is what I had), that is great. If you have to create a new domain, think about how it will reflect your new brand. Ideally, you want the same provider that manages your old domain to also manage your new domain to simplify the data migration issues. I had also set up a new Google Workspace account on the destination domain, which helped. All of this means understanding the various web, email, and other servers that you'll need, and what *Domain Name System* (DNS) records will support these services.

If you are going from old Google to new Google accounts, I will give you instructions as specific as I can. If you are moving to a different online neighborhood, you might have to make adjustments, or plan in advance how to get your content from A to B.

For example, I was using a *Google Voice* number that I have had for so many years it was still free of monthly subscription charges. Alas, if I needed to make any changes, I would have to set up a new Voice account with a new number and pay them a monthly fee for my new domain. One rule in Voice-land is that you can't transfer a (paid) telephone number unless you have owned the Voice account for at least three months[4]. It is these little "gotchas" that can ruin your migration schedule or uproot your plans. Best to have your plan all mapped out in advance, if you can.

Migration Time

Here is the workflow that I would recommend:

(1) Fixing your emails is a multiple-step process. Let's start with your password collection. Hopefully, you use a password manager, either a third-party tool or one built into your browser. I have a third-party tool for my password manager (I started with *Lastpass* and then moved over to *Zoho Vault*), and have collected more than 500 logins, most of which used my **`strom.com`** credentials as my user name. If you don't use a password manager, now would be a good time to pick one and start populating it as you go through your logins. Also, don't forget to change the login ID that you use to access the password manager's vault itself.

To do this process effectively, I first needed to pick a new email address to use. I didn't think it through when I began the migration process, however, and ultimately had to change some of my logins more than once. I am sure some talented programmer could automate this process, but in my case I had to individually sign on and see what I needed to do to change the account. In some cases, I already had created more than one account that used some other email address. The good news is that I pared down my list to "just" 350 logins. The bad news is that this process took a few weeks, because it was tedious and I wanted to make sure I did it correctly.

This step is the most time-consuming and most annoying, because email infrastructure is brittle, and relics of your old domain are kept in less obvious places. For example, two of my bank accounts used my **`strom.com`** email for notifications, but a made-up user name for logins. And as I was writing this article, I realized that I missed updating several other logins-with-custom-user IDs on my first pass through my password manager, and I had to make a more detailed inspection to find the other logins that buried the email notification settings. With one of them I needed to make a phone call (oh the horror!) to change my email address. So spend some time examining all your logins to ensure that you have removed all traces of your old domain while you still have it as an active address. For some of the logins, the site sent a confirming email to the old address (that required your acknowledgment) to ensure that the change was really from you. If you don't have access to the old address, it could cause problems.

Email addresses also lurk in various dark corners of your digital infrastructure for less obvious notifications. For example, LinkedIn has primary and secondary email addresses, and I forgot to change the latter one, and as I was writing this article I also remembered that I had to change the link to my website. Another example: I use the *WordFence* plugin to protect my WordPress website. There is a place deep within the WordPress menus where you have to specify where to send email alerts, as shown on the next page:

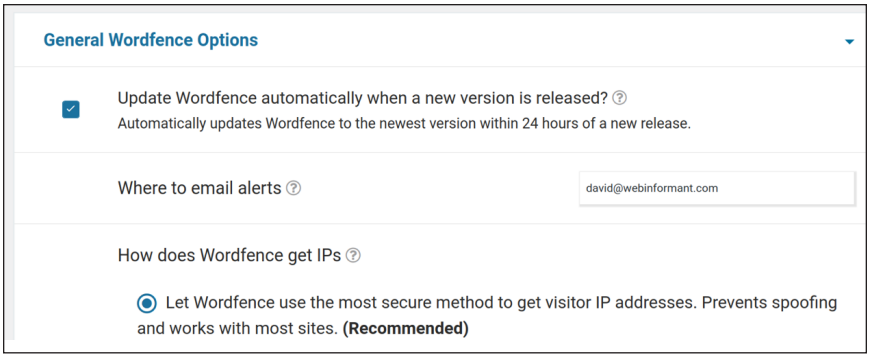

But email addresses are also used for mailing lists, and here again you will have to go on a scavenger hunt to seek out and change them. This process involves both lists that you are a member of (assuming you still want to get these messages) and lists that you create as list owner (such as using *Google Groups* or equivalent products). Now could be a good time to pare down your lists, especially if you find yourself deleting these messages more than reading them and can conquer the *Fear Of Missing Out* (FOMO) of not getting them in your future inbox. Some email list providers (such as Substack) make this transfer from old to new email easy, meaning you can transfer every subscription at once.

The previous discussion should convince you that you should make the transition from old to new email identity slowly, giving yourself and your systems time to catch these wayward notification messages. You will forget a few places, even if you are highly organized.

(2) Consider how to migrate your emails. Back when I first got online, I used a desktop email program. I think it was in 2005 or so that I finally moved over to Google's email services and have continued to this day. This consistency helps because there are ways to preserve your email archive, and I was able to migrate that archive to a new private Gmail address. You can do this in several ways:

- Use Google's own migration tool[5] (from the main Google Workspace Admin/Data Migration)—this tool is the newer of the various Google tools. I didn't try this way, but it seems to be the best way to migrate.
- Use Google's own import tool (All Settings/Accounts and Import), which will make a POP connection to your old account and gradually move everything over to your new account. In my case, when going from **`strom.com`** to a Gmail account, the program proceeded to fill my inbox with the thousands of decades-old emails that I long since had deleted and forgotten about. It did not migrate my labels, which I missed. Use this method at your own risk.

- Use other methods to migrate, such as *Google Takeout* or other third-party services. You would choose this method if you are migrating from *Microsoft 365* to Google, or vice-versa. You'll want to use a third-party transfer service, and there are numerous vendors. For these migrations, you might also want to create a local copy of your email archive (which is what Takeout does; more on that service in a moment). Going this route guarantees that not all the metadata (such as folders or names) will transfer authentically, and will require cleanup on the email aisle 7.
- You may want to not migrate anything, or minimally, and start afresh with a nice clean inbox and uncluttered account. If you use Google Workspace after it is disconnected from your old domain, it will remain intact, at least for the length of your subscription. I still had several months' worth of my account after the domain was transferred, and it was helpful to have the archive there as a backup, especially if I was searching for correspondence using my labelled emails.

If you use a Microsoft Exchange-based email, you have a much simpler migration process, basically swapping out your MX records from old to new domains. Craig Ellison, who has worked for various technology publications including *PC Magazine* and is now a retired technologist, hired a Microsoft VAR to do the actual work several years ago when he sold one of his domains. "I figured I was only going to do it once, so I outsourced it. I felt it was money well spent, and it freed my time for doing paid writing work. The moral is "do what you do best, and outsource the rest."

(3) Migrate your email contacts. I used *Google Contacts* with **`strom.com`** and wanted to keep using it with my new identity. That part was relatively easy after I found the right command location—you can export these contacts using the small icon on the top header line on the right-hand side, and import them in your new identity with the menu option on the left-hand side. If you are moving between Google and some other service, expect that not all of your data (groups, metadata) will transfer without some errors.

(4) Migrate your calendar appointments, if you have them as part of your domain. Google Calendar has both import and export options in the Settings page, along with the ability to create an offline calendar that you can use as a backup while you make the move. As with contacts, if you are moving between two services, expect errors in some data elements.

(5) Migrate video content on your YouTube account (if controlled by your domain). This method deserves special treatment, because YouTube has its own account ownership (called a Brand account) and verification methods. For example, you can add another account owner, but you have to wait a week before you can promote this person to "primary owner."

Google Takeout is great at making a complete backup copy of this content (along with other Google services that you may have forgotten about, such as my freebie Google Voice account). Takeout takes several days after you initiate the transfer—make sure you watch for the email notification (one of the options; it can also transfer data to Box and other online storage repositories).

(6) Don't forget about your documents and other online storage parts of your domain account. I had loads of documents, spreadsheets, and other digital effluvia. Again, Google Takeout is handy and will create a copy of all your files if you choose the right services to download.

(7) Moving your website(s) and other servers will depend on many factors. In my case, since I was moving from old to new domains that were both managed by my hosting provider (in my case, **`Pair.com`**), I edited my domain mappings in my hosting account, so the actual content didn't move anywhere, just the hosting provider's links to it. I had to edit some of my website HTML pages, and make a few phone calls to Pair to straighten some stuff out, and they handled that easily. I also had to add a plugin to my WordPress site (called Velvet Blue) to enable wholesale changes in the internal links from old to new domain names). I also needed to update the security certificates for the new domain name. If you use two providers, this step might involve a lot more coordination—and a lot more movement of the actual data that you wish to preserve.

(8) Next are the details about the transfer of the ownership of my old domain. For that domain, I was using GoDaddy as my registrar, and this transfer required unlocking the domain and specifying the transfer to my buyer's GoDaddy account. I was fortunate that we both used GoDaddy as our registrar, otherwise it would have necessitated me to work with **`Pair.com`**'s DNS servers.

(9) After you have migrated your content, it is time to finish the financial transaction. The typical method is to use an escrow service. They collect your money from the buyer and provide final control over the assets of your domain. Many businesses offer this service, and as I mentioned earlier, I used the aptly named **`escrow.com`** to finalize my sale. I also engaged a lawyer to vet the agreement that I signed with my buyer—something you probably want to do even if you have a broker representing you. You can specify who pays the escrow fee (which they are very transparent about) when you set up the sale on their platform. **`Escrow.com`** has three different transfer methods (buyer has new username and password, an authorization code is used to transfer between registrars, or in my case, a domain push from my GoDaddy account to the buyer's). I thought **`Escrow.com`** was first rate, and the transfer went smoothly.

(10) If you are still reading this piece, there is one more step: how to get the word out to your correspondents. I would recommend adding a line in your .signature block that mentions the new email address and domain details, and have some overlap between old and new accounts to give you time to tell everyone your new cyberspace coordinates as those emails roll into your old inbox.

Actually, there is yet another part of your workflow, and that is to check all your digital systems and services to ensure that you have done everything correctly. You will no doubt find a few logins that you missed in the initial inventory phase that have "hidden" embedded emails, or something that didn't quite rise to the surface that is now broken because of your new domain credentials. In my case, I got caught up in how to specify the right DNS settings so all my content would match their new cyberspace places. It also didn't help matters that I had to mediate among GoDaddy, Pair, and Cloudflare, either.

As you can see, migration of your domain and your digital assets is complicated, and it is made more so by the numerous special circumstances and oddball once-in-a-lifetime processes that you may or may not know your way around. Searching through Google support pages is a good first step at providing step-by-step instructions, and you can also find numerous YouTube instructional how-to videos if you get stuck. But the process will take more time than you anticipated, even for the most carefully planned workflows.

References and Further Reading

[1] David Strom, "So You Want to Sell Your IPv4 Address Block?" *The Internet Protocol Journal*, Volume 23, No. 2, September 2020.

[2] Three such pricing services include:
GoDaddy: **https://www.godaddy.com/domain-value-appraisal**
Atom: **https://www.atom.com/domain-appraisal/**
Domain Price Check: **https://pc.domains/basic/**

[3] Domain Broker: **https://www.vpn.com/domain-broker/**

[4] Transfer your Google Voice Number:
https://support.google.com/voice/answer/12083094?hl=en

[5] About migrating email with the new data migration service:
https://knowledge.workspace.google.com/admin/migrate/about-migrating-email-with-the-new-data-migration-service

[6] Move your YouTube channel from one Brand Account to another:
https://support.google.com/youtube/answer/3056283?hl=en

DAVID STROM has written several articles for IPJ, most recently on the history of the Interop ShowNet. He was the founding editor-in-chief for *Network Computing Magazine* (USA) , and he ran overall editorial operations for *Tom's Hardware.com*. He is the author of two books on computing, including one as co-author with Marshall T. Rose on the book *Internet Messaging: From the Desktop to the Enterprise* (1998, Prentice Hall). He lives in St. Louis and can be reached at his new email address: **david@webinformant.com**

Book Reviews

Code War

Code War: How Nations Hack, Spy, and Shape the Digital Battlefield, by Allie Mellen, Wiley, ISBN-13 978-1-394-28557-0, March 2026.

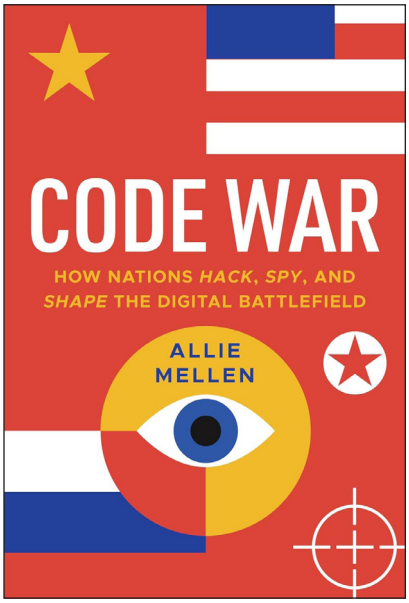

Allie Mellen has written an interesting book that takes the reader through a comprehensive historical narrative of the past several decades' worth of state-sponsored cyber attacks. While there have been numerous books on this topic, what makes this book unique is that she examines attacks that have been attributed to the US, Russia, and China, and shows their common and different approaches, and how they mix cyber warfare with their on-the-ground kinetic battles, such as what has and continues to happen in Ukraine over the past several years.

Mellen comes to this effort from a deep experience with cybersecurity, including five years as an analyst at Forrester Research and several jobs for private cybersecurity vendors.

Code War covers a lot of ground—from the earliest days of history to the present era, and how the modern digital age is just another way to repackage some of the ancient analog exploits. That deep historical coverage sets this book apart from other efforts that just skip lightly over the details and relevance of these antecedents.

Each country has separate ways that they approach cybersecurity, both from offensive and defensive positions. Each also has different contexts in which it evaluates its cyber efforts. The US context is to ensure its national security, maintain a strong economy, and support various freedoms. China wants to maintain its regime stability, protect its national interests, and regain control and influence in Asia. Russia wants to maintain economic stability, ensure its citizens are loyal to the regime, and remain a world superpower. These mixed goals compete and conflict with each other. And while it is great to have goals, the contradictions and conflicts among them make it hard for each regime to clearly evaluate and execute its cyber efforts.

Part of the problem, when seen in this tripartite context, is that the role and nature of the Internet is vastly different among the countries China's Internet is an instrument of state power, cultivated by absolute control. Russia's Internet is part of an hybrid digital/analog background of warfare against the world's democracies. And in the US, the Internet is part of maintaining a defensive and resilient digital ecosystem.

One element in common with these efforts is their work to isolate their residents from the global Internet community. These "splinternet" efforts restrict freedom of speech and as Mellen notes, it "becomes more difficult to spread democratic values globally." She chronicles the key steps of isolation and control of the Internet with a series of well-researched case studies.

Mellen proceeds to deconstruct operational playbooks of the three nations, and how they have used cyberattacks to fulfill their social contracts with their citizenry. The American chapters cover a wide range of cyber misdeeds, including one chapter that tells the stories about how Nathan Van Buren and Aaron Swartz independently ran afoul of various federal laws about computer network security. Swartz got caught illegally copying millions of academic research articles in his campaign to make this information more publicly available, eventually killing himself rather than cop a plea. Van Buren was a Georgia cop who was charged with illegal, unauthorized access to law enforcement databases, a case that went to the Supreme Court.

Another historical luminary is a story of how Ben Franklin constructed one of the first disinformation campaigns. Granted, the Internet was yet to be invented, but his playbook—using racist overtones—is very similar to many of the present day's digital campaigns. "Disinformation operations have always been part of the US experience, they are just more easily scalable with the Internet," she writes.

Another story concerns how in the mid-1800s, Edgar Allen Poe was part of an abysmal voting practice called *cooping*, whereby people voted early and often, receiving free booze for their efforts. Mellen uses this tale to take a closer look at how American voting practice has become more secure, despite exaggerated recent claims to the contrary. Such claims include the efforts of the Cybersecurity and Information Security Agency that was once a leader in securing our elections before it lost its mission, its director Chris Krebs, and at least a third of its staffers in 2024.

Most readers of *The Internet Protocol Journal* are familiar with the stories about how Iran and Russia hacked our 2016 and 2020 elections, but Mellen dives into the details, showing how Iran, for example, tried to alter the final voting tabulations in 2020. Also a familiar tale for many readers is the plight of Phil Zimmerman, inventor of *Pretty Good Privacy* (PGP), and how it became a legal lightning rod and the first technology to be designated a war-based munition. This has echoes of the current day whereby the Defense Department can designate Anthropic's AI similarly (and perhaps equally unjustly).

Most of us are familiar with China's *Great Firewall*, but Mellen describes its companion isolation and protective programs, including the *Golden Card Project* (its own online financial network) and the *Golden Shield Project* (its national surveillance and censorship network). Some of these containment efforts have been abject failures, such as the *Green Dam* software that was a required application begun in mid-2009 to be installed on all Chinese computers and phones. The software was buggy and so unwieldy that the state gave up the project within a few months.

Mellen analyzes numerous Russian attacks and describes four common elements of their playbooks:

- Denial of service attacks, including *Global Positioning Service* (GPS) and satellite jamming,
- Traditional espionage operations,
- Psychological operations, such as phishing, disinformation, and audio/video deepfakes, and
- Malware-based data wipers.

Each of these elements has evolved over time, and carries its own hybrid physical attack vectors to amplify the attack. As I mentioned earlier, Ukraine is where all four of these elements were brought together alongside the physical warmaking machinery to form a single continuous battlespace.

Mellen's tour through history and technology shows how political leadership has failed to live up to promises with its citizenry to maintain and improve their respective social contracts: China's prosperity is crumbling, Russia's safety is evaporating, and America's economic divide continues to worsen. By having this deep historical dive, the reader can see where things went off the rails, and why.

Missing from her excellent treatment of world powers is a focus on Iran, although it is mentioned briefly in several case studies. Also missing is more than a passing glance at AI.

Mellen concludes with a dark vision of the "fourth power," that of the major tech companies who treat their users as "digital peasants living in a world of corporate feudalism. Users till the soil (creating data), pay taxes (such as subscription fees), and live in castles (the digital platforms themselves), having no say in how the kingdom is governed." The real nation states like China, Russia, and the US and the digital nation-states such as Google, Apple, and Meta all want your data and your attention so they can exploit you and leverage your resources.

—David Strom
david@webinformant.com

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**

Independence in Cyberspace

Declaring Independence in Cyberspace: Internet Self-Governance and the End of US Control of ICANN, by Milton L. Mueller, ISBN-13 978-026255258, MIT Press, 2025.

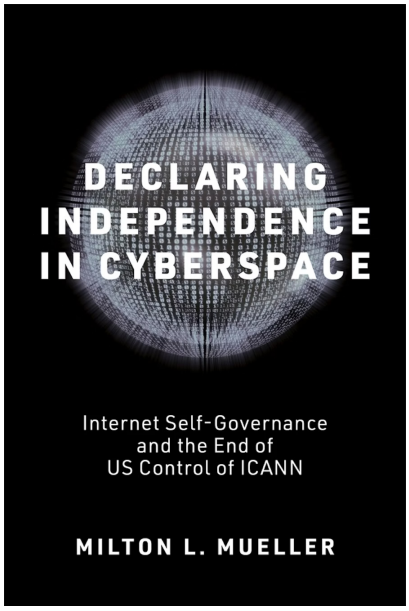

Milton Mueller's latest book, *Declaring Independence in Cyberspace*, claims to be an "independent scholarly account and evaluation of ICANN." Mueller is a professor at The Georgia Institute of Technology School of Public Policy and one of the main organizers of the *Internet Governance Project* that is associated with that institution. This book is his latest work, following a series of other books he has written on earlier periods of Internet history.[1,2]

While his scholarship is very much in evidence with the long list of citations and pages of links to various reference documents, I am not sure it succeeds in being a truly independent evaluation, or a book worth reading from cover to cover. Clearly Mueller "has an axe to grind" and a point of view to espouse in his tale of the evolution of the *Internet Corporation for Assigned Names and Numbers* (ICANN) from a private company paid for by the US government to what it is today. He calls this transition an *experiment*, which if taken literally, means we are all lab subjects that are still unwitting participants. That attitude infects and detracts from his story.

The book is concerned with the ICANN transition years beginning in 2014. It covers mostly the same period that is documented in one part of *Geopolitics at the Internet's Core*, which I reviewed in the December 2025 issue of IPJ.[3,4] Mueller's turgid academic writing style makes it hard to parse and place in context his dense narrative of events on the "nooks and crannies" of Internet Governance. He is fond of highlighting many of the debates in which he participated as a vociferous commenter. That may be fine if *Declaring Independence in Cyberspace* were a reference work, but it falls short as a compelling story. It is more a historical reference work.

Mueller's book starts with the politics and processes and then moves inward to show the slow, steady freight train of progress held in various meeting rooms and policy papers and consensus building (or lack thereof) exercises. Because he self-characterizes his role as "ICANN Official Nemesis," in his LinkedIn profile[5], it is hard to separate the history from his commentary and his motivation to point out the numerous failures of nearly everyone else involved in its transition.

Nevertheless, there are some particularly useful places worth reading. I would recommend starting his book with chapter 14, rather than at the beginning. This chapter deals with the post-transition of ICANN and how it has evolved since freeing itself from the shackles of the US Department of Commerce.

For example, he documents three post-transition challenges, including the WHOIS battle brought about as a result of European privacy regulations, attempts to sell off the **.ORG** registrar, and the attempt to delete the Russian **.RU** country code as a result of the Ukraine war. The descriptions of the situation and actions by the various stakeholders in resolving these conflicts is perhaps the best written explanation of these three events that I have seen in one place. My reason is that there is a ready plot line through each story, with a beginning, a middle, and a resolution. If only the remainder of his book took this tactic.

The remainder of his book goes into excruciating detail about his view of *multistakeholderism*: how he defines it, how ICANN and the Internet Registries use or abuse it, and his recommendations on what should be the ideal policy collection. He is so caught up in his own point of view that it is difficult to see how his recommendations could be enacted by any policy-maker. The last chapter posits three future end-states for ICANN: He starts out by mentioning its current path, that of survival and right-sizing its policy-making process in the face of creating thousands of new *Top Level Domains* (TLDs). After all the sniping about what ICANN should have done, he admits that the "transition accomplished what it intended." Yet he still presents specious arguments here, such as suggesting that "the market for domain name registrations stops growing or shrinks."

The other two potential ICANN futures—a total takeover by the United Nations which results in handing control over to various national state actors, or a scenario where the US and China split the Internet into their own walled tech gardens complete with separate and conflicting standards and protocols—aren't very palatable alternatives, I fear.

The two *Declaring Independence* and *Geopolitics* books are vastly different in tone, topics, and intended audiences. *Declaring Independence* is for a more limited audience and has less appeal to many IPJ readers, because it doesn't have the underlying technical meat that is part and parcel to *Geopolitics*. The *Geopolitics* book is a lot more readable and engaging, with more insider viewpoints into what was enacted (versus what is debated and seen under Mueller's microscope), and approaches the transition as a prelude to what ICANN is doing today in a more thematic and relatable story. As I mentioned in my review of *Geopolitics*, you can learn a lot more than who-said-what-at-which-meeting, which seems to be Mueller's goal with his book. *Geopolitics* is a protocol-first book, explaining the underlying protocols as a proxy for explaining the ultimate politics, policies, and processes involved in running the Internet. *Declaring Independence* goes into more external details, such as a long section on how Edward Snowden's revelations came at a critical time in Internet Governance and other aspects such as the creation and controversy over the **.XXX** domain—both of which were not even mentioned in passing in *Geopolitics*.

Fans of Mueller's earlier works and copious blog posts will welcome Declaring Independence. But I fear this is an inherently small audience.

—David Strom
david@webinformant.com

References

[1] Dave Crocker, "Book Review: Ruling the Root," *The Internet Protocol Journal*, Volume 5, No 4, December 2002.

[2] Milton Mueller, "Letter to the Editor, regarding Ruling the Root," *The Internet Protocol Journal*, Volume 6, No. 1, March 2003.

[3] *Geopolitics at the Internet's Core*, by Fiona M. Alexander, Laura DeNardis, Nanette S. Levinson, and Francesca Musiani, ISBN 978-3031894770, Springer Nature, 2025.

[4] David Strom, "Book Review: Geopolitics at the Internet's Core," *The Internet Protocol Journal*, Volume 28, No. 3, December 2025.

[5] Milton Mueller LinkedIn profile:
https://www.linkedin.com/in/miltonmueller/

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.

In Memoriam: David Jack Farber (April 17, 1934 – February 7, 2026)

ISC recognizes "Grandfather of the Internet"

The *Internet Systems Consortium* (ISC) and the entire Internet community are mourning the loss of our friend and colleague, Dave Farber. Farber was one of the original members of the ISC Board of Directors, appointed at ISC's incorporation in 1994.

Dave was a professor of computer science, noted for his major contributions to programming languages and computer networking.

Dr. Farber graduated from the Stevens Institute of Technology in 1956 and began an 11-year career at Bell Laboratories, where he helped design the first *Electronic Switching System* (ESS-1) and the *SNOBOL* programming languages. He subsequently held industry positions at the Rand Corporation and Scientific Data Systems, followed by academic positions at the University of California, Irvine, and the University of Delaware.

At Irvine his research work was focused on creating the world's first operational distributed computer system. While a member of the Electrical Engineering Department of the University of Delaware, he helped conceive and organize the major American research networks; the *Computer Science Network* (CSNET), the *National Science Foundation Network* (NSFNET), and the *National Research and Education Network* (NREN). He helped create the NSF/DARPA-funded *Gigabit Network Testbed Initiative* and served as the Chairman of the *Gigabit Testbed Coordinating Committee*.

Dave subsequently was appointed Alfred Fitler Moore Professor of Telecommunication Systems at the University of Pennsylvania, where he also held appointments as Professor of Business and Public Policy at the Wharton School of Business and as a Faculty Associate of the Annenberg School for Communication. He served as Chief Technologist at the US *Federal Communications Commission* (2000–2001) while on leave from the university.

He also was a Distinguished Career Professor of Computer Science and Public Policy at the School of Computer Science, Heinz College, and Department of Engineering and Public Policy at Carnegie Mellon University. He was a Fellow of the IEEE, ACM, and the AAAS.

In 2018, at the age of 83, Dave moved to Japan to become Distinguished Professor at Keio University and Co-Director of the Keio *Cyber Civilization Research Center* (CCRC). He loved teaching, and taught his final class on January 22, 2026.

Dave was a delightful person, part of the bedrock of the Internet, and a great friend to ISC over the course of decades of board membership. We will miss him. We extend our deepest condolences to Dave's family and friends.

—Jeff Osborn

A Tribute to Professor David J. Farber

On Saturday, February 7, 2026, our great mentor and dear friend, Professor David J. Farber, passed away at the age of 91.

Throughout his illustrious career, Dave nurtured countless leaders of the computer science and Internet communities in the United States. While I consider myself but one of his humblest students, I was incredibly fortunate to have known him since the dawn of computer networking in Japan. From our early days of guidance to our collaboration on trans-Pacific connectivity, and finally, to the realization of my long-held dream of inviting him to Keio University—where we served together as Co-Directors of the *Cyber Civilization Research Center* (CCRC) at *The Keio University Global Research Institute* (KGRI)—the deep bond we shared was a source of unexpected and profound happiness in my life.

Our journey began when Dave visited Japan alongside Mark Horton of *UUNET Communications Services* (UUNET) and Professor Larry Landweber of the Computer Science Network (CSNET), just as we were launching *Japan University NETwork* (JUNET). Even before then, Dave had maintained close friendships with my seniors, Professor Hiroshi Inose of the University of Tokyo and Professor Hideo Aiso of Keio University. It was through these connections that he first took me under his wing with such kindness.

In the early days of interconnecting with CSNET, Dave and I worked together from the University of Tokyo to run MMDF for email exchange over X.25. Later, when we implemented *Serial Line IP* (SLIP) with Professor Hideyuki Tokuda of *Carnegie Mellon University* (CMU), we faced the constant thrill of technical uncertainty. In those days, a simple mail loop could lead to exorbitant packet-based charges, forcing us into the "thrilling" routine of somehow justifying telecommunication costs to our universities.

Yet, despite the risks, we were driven by a singular spirit: "We want to connect the world." The researchers of that era were truly united by this mission. Dave and I were, in every sense, brothers-in-arms.

The challenge of international communication costs was eventually resolved thanks to another of Dave's close friends, Professor Inose. He appointed me, then a research assistant at the University of Tokyo, to lead the email component of an *National Science Foundation* (NSF) project providing access to Japanese databases via leased lines.

I proposed to Professor Inose—a staunch advocate for *Open Systems Interconnection* (OSI) and international standard protocols—that we use X.25 bandwidth where "terminals and email could coexist." By running IP over X.25, we successfully interconnected the WIDE domestic IP network with the CSNET IP network. This became the foundation for the stable interconnection of the Internet in Japan.

By the late 1980s, as the WIDE Project was gaining momentum, I was still in my 30s and holding the title of research assistant. Breaking through the thick walls of regulation and bureaucracy to connect Japan to the world was an immense struggle. It was then that I made a proposal to Dave, Larry, and Vint Cerf: "In the U.S., you may be understood, but in Japan, we have struggled tremendously just to allow universities to connect freely overseas. To ensure that not only Japan, but all of Asia and the world, can interconnect, can we not create an organization with the necessary authority?"

Their response was immediate: "Jun, we will create a corporation called the *Internet Society* (ISOC). Will that work?"

Thus, in 1991, the Internet Society was born. In appreciation of their efforts, the WIDE Project hosted INET'92 in Kobe—the first ISOC international conference of its kind. Throughout the birth of the Japanese Internet, Dave's powerful support and encouragement were always there.

I know that the reason Dave loved Japan and stayed by our side was not solely due to our friendship. He also had a profound love for Japanese cuisine. We unashamedly made use of this fact, patiently inviting him until we finally welcomed him to Keio University in 2018 as a Distinguished Professor and Co-Director of the *Cyber Civilization Research Center* (CCRC).

Since then, he mentored three or four new generations of students with genuine warmth and closeness. As a place dedicated to studying the entirely new civilization born from technology, I believe the name "Cyber Civilization Research Center" was a fitting home for Dave's final years of work.

As this new civilization continues to evolve, we will never forget our gratitude to Dave Farber, nor will we forget his passion, which remained undiminished until his very last day.

Dave, please continue to watch over us as we, who have inherited your wisdom, forge the path into the future.

Thank you, Dave. Rest in peace.

—Jun Murai
Co-Chair, Cyber Civilization Research Center

References

[1] *Internet Hall of Fame*, 2013 Inductee: Dave Farber.

[2] Simson L. Garfinkel and Eugene H. Spafford, "In Memoriam: David J. Farber," *Communications of the ACM*, Volume 69, No. 4, April 2026.

[3] Peter Wayner, "David J. Farber, 'Grandfather of the Internet,' Dies at 91," *The New York Times*, February 14, 2026.

Fragments

WSIS+20 Outcome Document Published

On December 17th, 2025 the United Nations published "Outcome document of the High-Level Meeting of the General Assembly on the Overall Review of the Implementation of the Outcomes of the World Summit on the Information Society."[1] According to the Internet Society[2,3], the important aspects of this document are:

- It makes the *Internet Governance Forum* (IGF) a permanent forum of the United Nations. By securing a lasting mandate, the IGF community is now ready to plan for the long term, including sustainable financing and a clear program path. The Internet Society and its community stand ready to help shape that next chapter.
- It reaffirms the importance of the WSIS framework, including the Tunis Agenda's working definition of Internet governance—which underpins the multistakeholder model—and the WSIS Action Lines—which turn that vision into a practical implementation program.
- It establishes a "joint implementation roadmap" to align the WSIS+20 follow-up with the *Global Digital Compact* (GDC), ensuring that future governance efforts remain unified rather than fragmented.
- It recognizes the importance of the global network of National, Regional, and Youth IGF initiatives, validating the grassroots work of our community.

References and Further Reading

[1] UN WISIS+20 Outcome Document, available in the six official UN languages: **`https://digitallibrary.un.org/record/4095872?ln=en&v=pdf`**

[2] Israel Rosas, "WSIS+20 Reaffirms Multistakeholder Governance and a Lasting IGF," *Internet Society Blog*, December 19, 2025.

[3] Carl Gahnberg, "From Commitments to Practice: Internet Society's Priorities for WSIS+20 Implementation," *Internet Society Blog*, April 13, 2026.

[4] Geoff Huston, "Opinion: The End of Multi-Stakeholderism?" The *Internet Protocol Journal*, Volume 28, No. 3, December 2025.

[5] Avri Doria, "Counter-Opinion: Multi-Stakeholderism Is Not Ending; Rather, It Is Moving to a Next Stage," *The Internet Protocol Journal*, Volume 28, No. 3, December 2025.

[6] Avri Doria, Postscript at the end of: "m17m is not ending, but the honeymoon is over," *Medium Blog Post*, January 2, 2026.

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
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
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