The Internet Protocol Journal

May 2025 Volume 28, 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

The *TCP/IP Interoperability Conference*—later renamed *Interop*—began as a small workshop in August 1986. It quickly grew in scope to incorporate tutorials, and by 1988 an exhibition network connected 51 exhibitors to each other and to the global Internet. This network was designed and deployed by a group of volunteers, and it became the proving ground for many emerging technologies. In 1994, Interop added Tokyo to its international venues, where 30 years later the conference and exhibition attracts more than 120,000 visitors annually. Following an article by David Strom describing the history and evolution of the Interop show network in our previous issue, we now bring you the first installment of an article that describes how this network continues to evolve at the Tokyo Interop show. The article is by Takashi Tomine, Ryo Nakamura, and Ryota Motobayashi—all members of the team that designs and deploys their version called *ShowNet*. A second installment detailing the technologies demonstrated in the 2024 ShowNet will be published in a future issue.

Our previous issue also contained an article about the "IPv6 Mostly" experiment that was conducted during APRICOT 2024 in Bangkok. It is perhaps surprising that we are still very much living in an Internet that is heavily dependent on *IP Version 4* (IPv4) given the amount of time that has passed since the initial *IP Version 6* (IPv6) specifications were published. In our second article, Geoff Huston provides an in-depth analysis on the topic of IPv6 Transition and suggests that perhaps changes in Internet Architecture and technological developments will have us waiting a very long time before IPv4 addressing becomes obsolete.

Book reviews used to be a fairly regular feature in this journal, but it has been quite a long time since we have published any reviews. We asked Craig Partridge to review the book *The Real Internet Architecture: Past, Present, and Future Evolution,* and we hope this latest review will encourage you to send us suggestions for other books on networking and related topics. As always, you can contact us with your feedback by sending an e-mail to: **ipj@protocoljournal.org**

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

ShowNet at Interop Tokyo:
A Continuously Evolving Demonstration Network

by Takashi Tomine, Ryo Nakamura, and Ryota Motobayashi

Interop Tokyo, which inherits the objectives of the Interop conference series, is the largest annual exhibition of Internet technologies in Japan. It is held yearly for three days in June. Over 500 exhibitors showcase their products and services at individual booths, and more than 120,000 people visit the venue during the exhibition, as shown in Figure 1. Moreover, a co-located conference offers several dozen sessions and keynote talks.

Figure 1: A view of the Interop exhibition in 2024.

An essential part of Interop Tokyo is *ShowNet*, the largest demonstration network built at Interop Tokyo exhibitions. ShowNet provides network connectivity for Interop exhibitors and attendees but is not limited to this service. Since Interop originates from the word "interoperability," ShowNet conducts various interoperability tests, experiments, and demonstrations of new networking technologies. For example, in 2019, we deployed service function chaining using *Segment Routing over* IP*v6* (SRv6)[1] with four SRv6-capable nodes and five SRv6 proxies[2] from different vendors.

At that time, SRv6 was an emerging packet-forwarding paradigm, and we faced varied challenges and considerations to achieve this archetype while serving user traffic.

The knowledge we gained through the deployment was published as an Internet Draft[3]. We have deployed and demonstrated not only routing techniques but also broader technologies, including facilities, optical transport, wireless, security, monitoring, testers, and emerging technologies, such as 5G and media over IP in recent years.

The 2024 ShowNet, featuring comprehensive technical demonstrations, consisted of more than 20 full-height racks, allowing attendees to see the devices running in production. Figure 2 shows a picture of such a ShowNet booth in 2024.

Figure 2: The ShowNet booth in 2024. Whiteboards were mounted on the side walls of each rack, where the NOC members wrote explanations about the devices, design, and technologies. Attendees could see the running devices with those explanations.

History

Interop Tokyo celebrated its 30th event in 2023. In other words, thirty years have passed since an IT trade show, *NetWorld+Interop*, landed in Japan with cutting-edge technologies such as Ethernet 10BASE-T, *Fiber Distributed Data Interface* (FDDI), *Asynchronous Transfer Mode* (ATM) with IP, *Xerox Internetwork Packet Exchange* (IPX), and Apple's *AppleTalk*. Because of a consolidation among other strong Informa sister brands, no Interop event has been held in the United States since 2023. However, Interop is still alive and well in Tokyo, and it maintains its original mission: establishment of multi-vendor interoperability.

In the early 1990s, fathers of the Internet in Japan who visited an Interop event in the US were impressed by its effectiveness—a practical display of interoperability among multi-vendor networking equipment. The groundwork to adapt the event to a Japanese audience began then, and in 1994 Tokyo became one of the host cities of *NetWorld+Interop*, with Las Vegas, Berlin, Atlanta, and Paris (Interop was merged with Novell's *NetWorld*, a similar event that occurred from 1994 to 2004).

As an essential part of the event, Interop ordinarily deploys a temporary show-floor network *InteropNet*—or *ShowNet*, varied by years or venues—to provide Internet connectivity for each exhibitor. This concept was naturally introduced for the Tokyo show too. Until 1997, the design and equipment of InteropNet was basically shared through every show during the annual world tour. Actually, for the first Tokyo show, the construction and verification work of InteropNet (for Tokyo) was held not in Japan but at the Ziff-Davis's *Hot Stage Test Facility* in Sunnyvale, California. The persons of talent for the latest network construction and operation—the *Network Operations Center* (NOC) team—were also invited globally during the initial Tokyo era.

In 1998, following the event organizer's business operations review, Tokyo decided to set up a show-floor network with a local focus. The Tokyo NOC team has focused solely on the Japanese events since that time.

After making that decision in 1998, Tokyo has refrained from using the original *InteropNet* name, and now calls its own network *ShowNet*. ShowNet has since run every year except 2020, when the COVID-19 pandemic prohibited such public gatherings. In addition to Japan's Internet technology community, diverse members from industry and academia now gather every year to continue building and demonstrating ShowNet.

Volunteers are indispensable to achieving such complex networking. Initially named *InteropNet Team Members* (ITMs), volunteers are currently called *ShowNet Team Members* (STMs) in Tokyo. This program, which includes an educational aspect for young students and engineers, continues to be an essential component of ShowNet.

Who Makes ShowNet?

Building ShowNet at the Interop Tokyo exhibition is not an easy feat, so more than 650 engineers with diverse backgrounds are now involved in the project. The NOC team includes the core members, who design the ShowNet network and conduct broad experiments and demonstrations. In recent years, the NOC team has consisted of around 30 expert volunteers from academia, carriers, vendors, etc. They use their areas of expertise and skills to manage the project. Two leaders supervise the ShowNet project; they choose the NOC team members yearly. Teams are either selected by invitation from personal and professional connections or transitioned from STMs or Contributor Members.

The STM program offers a unique opportunity for university students and junior staff from companies to obtain hands-on experience in network operations. Participants in the STM program, who are relatively young and novice network engineers, engage in building ShowNet at the venue as volunteers, and they have the opportunity to touch, configure, operate, and debug various devices and cutting-edge technologies.

This valuable experience is difficult to gain in universities or in their regular workplace. The program also allows young engineers to build and foster relationships and learn from each other by spending two weeks building ShowNet together. In recent years, around 30 slots have been available for participants in this program, but we receive more than twice as many applications each year, so the NOC team members are responsible for the selection process. Figure 3 shows the 2024 STMs.

Figure 3:
The participants in the STM program in 2024.

ShowNet Team Members engaged in building ShowNet.

The third category of engineers involved in ShowNet is the *Contributor Members* who showcase products at Interop Tokyo each year. The contributor vendors make their products and services available to ShowNet and demonstrate them on the live network during the exhibition. These members are skilled engineers from those vendors, and they help build ShowNet with their expertise in the products. The presence of the contributor members is also indispensable for building ShowNet.

A Timeline in a Year

This section briefly introduces a timeline of ShowNet in a year. We, the people involved in ShowNet, put a long-term effort into accomplishing the ShowNet project every year.

Planning

ShowNet covers broad aspects of today's networking technologies. To manage this complexity, we organize the project into working groups, each focusing on a specific field. In 2024, we had 11 working groups leading the following fields: facilities, optical transport, external connectivity, backbone network, data center and cloud, wireless network, monitoring, security, testers, 5G, and media over IP. The NOC team consisted of approximately 30 members, with each working group led by two to four NOC team members who have expertise in that group's area of interest.

Preparation for ShowNet starts in October, the year before the Interop exhibition in June. First, the two leaders gather and organize the NOC team members and begin to discuss topics and technologies that ShowNet will address in the next Interop Tokyo. Then the leaders meet monthly with all NOC team members to share and discuss the overall structure and design of demonstrations. Additionally, each working group holds meetings at least once a month, as needed.

The Contributor Members—vendors providing their products to ShowNet—join the discussion in December. The NOC team members introduce the concept for the next ShowNet and the technologies they want to adopt to the contributor members per working group. Also, the contributors propose their products and use cases they wish to showcase. The NOC team members receive these requests and integrate them into demonstrations. From then until the end of May, the demonstration contents are continuously refined, and the NOC team members consolidate everything into a concrete network design.

Hot Stage

Two weeks before the Interop Tokyo exhibition, we start building ShowNet at the Makuhari Messe exhibition hall. Building ShowNet has two phases: *Hot Stage* and *Deployment*. During hot stage, we build and test all the designs and conduct planned interoperability tests and experiments. In recent years, we have allotted eight days for the hot stage, and we hold two all-hands meetings every day, one in the morning and one in the evening, to share progress as we continue the construction of ShowNet.

When the hot stage begins, all members of the NOC team, the STM, and the contributor members gather at the venue and start building the network. First, we install every device in the right place on the racks, turn on the devices, and check their status. Checking device status is very important because some devices are transported directly from overseas to the venue, so it is necessary to ensure that they are not malfunctioning. We usually finish this process on the first day.

On the second day, we start the network setup: connect appropriate links between devices with patch cables as designed. After the physical network connections are completed, NOC members in charge of the backbone network start configuring the backbone routers with the help of the ShowNet team members. In the early days, ShowNet backbone was a simple Layer-3 network with a single *Interior Gateway Protocol* (IGP) instance. But now, ShowNet adopts several overlay technologies such as *Multiprotocol Label Switching* (MPLS), SRv6, and *Virtual eXtensible Local Area Network* (VXLAN), so we have to configure more overlays after the Layer-3 routing configuration.

The working groups other than the backbone network group prepare their demonstrations in parallel. Every part of ShowNet is built with multi-vendor equipment, so we have to check interoperability everywhere. The working groups also conduct several interoperability tests during the hot stage. These interoperability tests are beneficial for finding bugs or slight differences in implementation.

Sometimes, these bugs or differences are critical to building ShowNet, so contributor members from vendors try to fix them with their development teams.

Testing the network is always essential, even in an event network. In ShowNet, we conduct failover tests in the latter part of the hot stage—stop and resume each backbone router sequentially and confirm that routing redundancy works as expected. If troubles arise during the test, the backbone network group of the NOC team and the ShowNet team members troubleshoot and debug the problems together. This collaborative troubleshooting process is also a good hands-on experience for the junior network engineers of STM. Figure 4 shows a snapshot of a failover test.

Figure 4: A snapshot of the failover test in 2024. Red lines in the display indicate that some user segments have unexpectedly lost the connectivity, and the NOC team and ShowNet team members start to troubleshoot. The software used here is deadman[5], which was designed and implemented for ShowNet.

Deployment

In 2024, after we finished the hot stage, we started to deploy the ShowNet network in the whole Makuhari Messe venue four days before the Interop Tokyo exhibition began. Interop Tokyo used five halls in Makuhari Messe. The ShowNet network spread to each hall with optical transport from the ShowNet booth. Every hall had a small booth on which access switches were installed for ShowNet to extend the network to all the exhibitors' booths. Electrical construction members deployed optical fibers from the ShowNet booth to the small booths in each hall and copper cables from the switches on the booths to the exhibitors' booths. After spreading the cables, the ShowNet team members connected these cables to the access switches of the ShowNet backbone and checked the correctness of Layer-1 to Layer-7 connectivity. Figure 5 shows a snapshot of such a scene.

Figure 5: Three ShowNet Team members and a NOC team member are checking deployed cables for exhibitor booths in an exhibition hall.

As the deployment phase begins, exhibitors of Interop Tokyo also arrive at the venue and start preparing their booths. ShowNet provides Internet connectivity for the demonstrations held in their booths. If they have any problems on the ShowNet network, they go to the *ShowNet Service Counter* and describe their problem, and we start to identify and resolve it. Usually, problems in this phase are caused by physical things like a cut cable or mis-connection of cables or access switches, but sometimes some logical bugs cause critical challenges. Such bugs are sometimes difficult to solve because we must fix them before the exhibition starts.

It is also crucial to ensure the visibility and presentation of the equipment and services contributed to ShowNet. After the whole ShowNet network is built, we tidy up the ShowNet booth. We try to ensure that every piece of equipment inside the racks is presented well, because it is not only a device but also an exhibit. We also post captions for all equipment and prepare description slides for attendees.

After we finish all the processes for building the ShowNet, we complete the network diagram. Figure 6 shows the network diagram of ShowNet 2024. You can see all the devices, links, services, and designs of ShowNet 2024 on this diagram. One of the NOC members creates the diagram. From the hot stage onward, the same member continuously monitors all the design and configuration changes and keeps the diagram up-to-date. Eventually, the diagram captures all of the ShowNet network on a single sheet. This diagram is an essential tool for ShowNet: engineers use it to grasp the overall network design, communicate and share changes, and troubleshoot problems. The diagram is practical and functional, especially when building the network and troubleshooting.

Figure 6: Diagram of the ShowNet network at Interop Tokyo 2024. The full-size version is available[6]. Green icons indicate Layer-2 devices, red icons are Layer-3, and yellow icons are Layer-4 and above devices. The icons in the diagram are also available under the Creative Commons license[7].

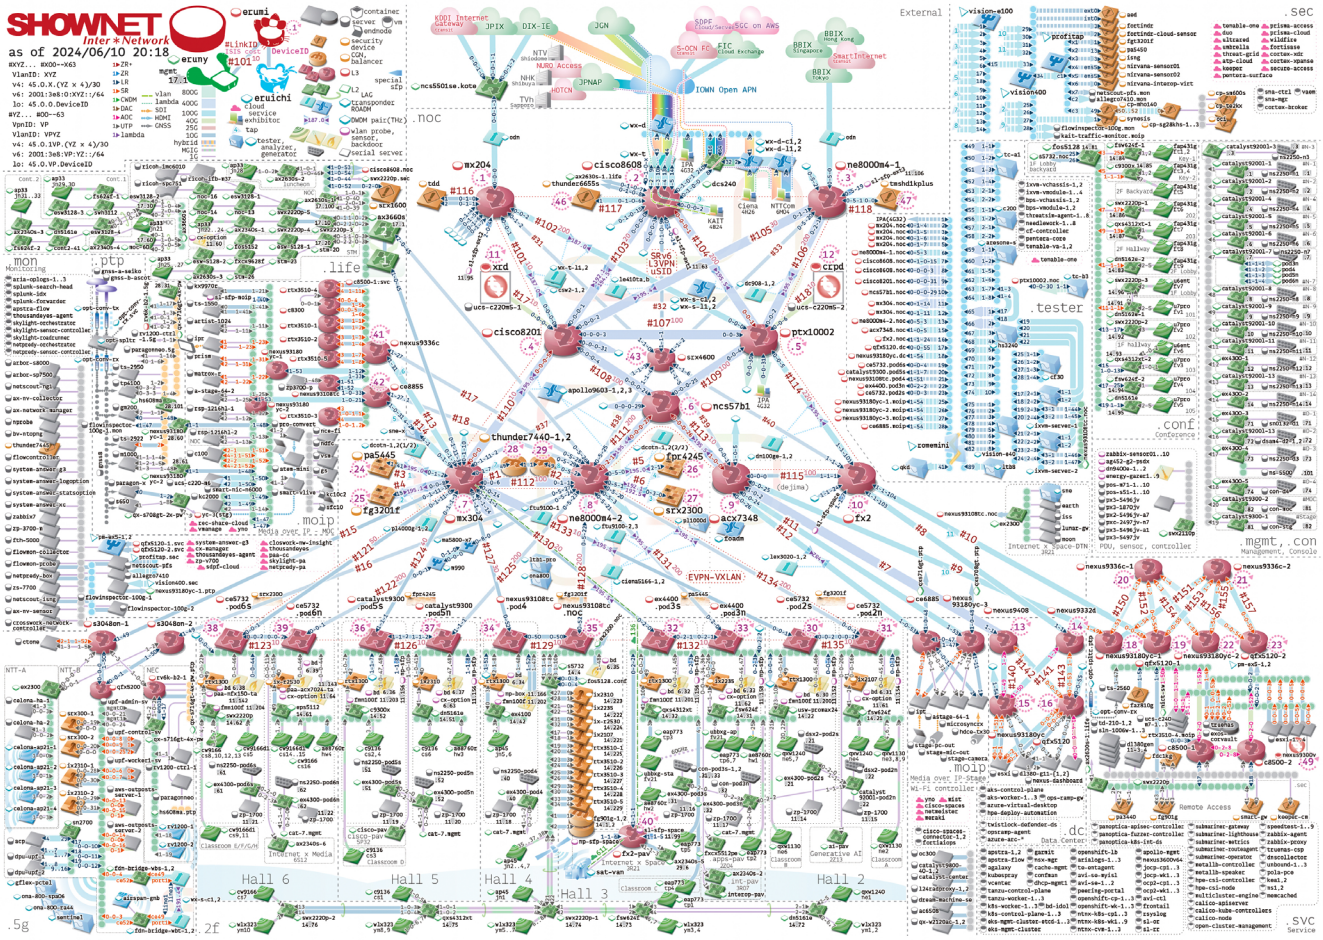

Exhibition

The Interop Tokyo exhibition usually starts on a Wednesday in mid-June and lasts three days. When it starts, we begin operating the ShowNet network and also offer some presentations for attendees. ShowNet is one of the main parts of Interop Tokyo, so many attendees come to the ShowNet booth. Attendees can see the demonstrations, learn about the technologies, observe the devices running in the racks, and see and feel the functions that run on ShowNet.

During the exhibition, ShowNet members, especially the NOC team members, work on showcasing the ShowNet network. The NOC room inside the ShowNet booth is equipped with large screens that show several tools monitoring the network status and security, as shown in Figure 7. The NOC team members operate ShowNet from there, and attendees can observe them; actually, the NOC room is also a part of the show. In addition, the NOC team members conduct tours, called *ShowNet Walking Tours*, during the exhibition. Participants on the tours can walk around each ShowNet rack with a NOC member and observe the running devices, and the NOC member will answer any questions. Figure 8 captures a scene from one of the tours. The NOC team members also give several presentations at the ShowNet booth and in session rooms at the exhibition.

Figure 7: The NOC room on the exhibition floor, visible to attendees.

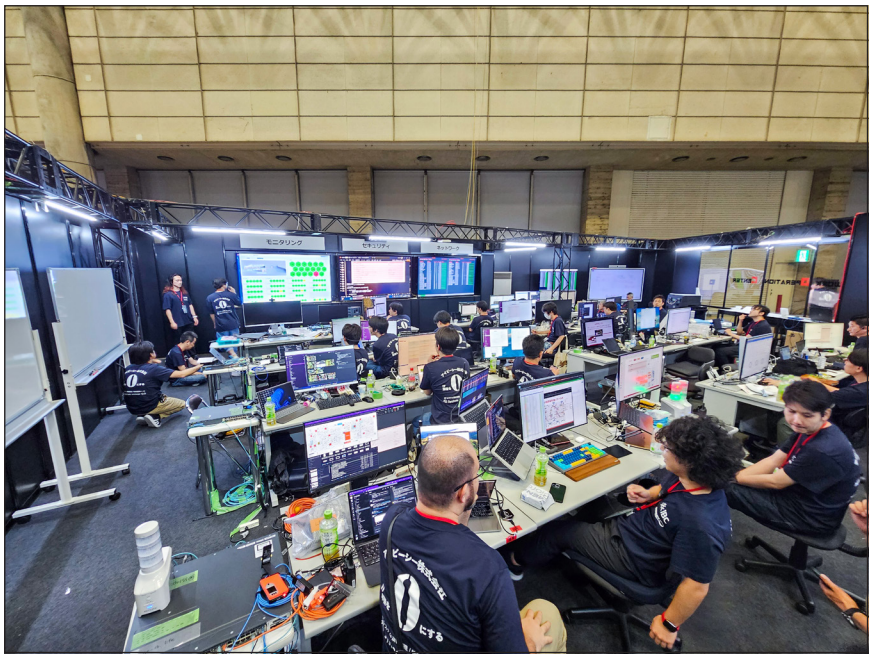

Tear-Down

After the 2024 exhibition ended at 5 p.m. on Friday, we started to tear down the ShowNet network. Our contract required that we vacate the halls by midnight. First, the NOC team and contributor members shut down devices, a requirement before they could be powered off. Next, we shut off all power supplies, unplugged the patch cables, unmounted the devices, and returned them to the contributors. After that, the NOC and ShowNet team members wound up all cables and cleaned up all racks for use next year. Figure 9 shows the racks during tear-down.

Conclusion

The ShowNet network is a unique environment. It not only provides Internet connectivity for exhibitors and attendees, it also displays a large-scale ephemeral event network that will demonstrate emerging and cutting-edge technologies. ShowNet conducts various interoperability tests, experiments, and demonstrations with numerous devices contributed by multiple vendors. Furthermore, ShowNet offers an invaluable opportunity for engineers to collaborate with diverse engineers from different fields. We believe that the connections and relationships among them established through ShowNet have contributed to revitalizing network communities. In addition, it has more than 30 years of experience and adds to our knowledge to handle these cutting-edge trials. Interop Tokyo will continue this work inherited from US Interop.

Acknowledgments

We would like to express our gratitude to everyone around the world who has been involved in all Interop events, from the past to the present.

Figure 8: ShowNet Walking Tour: NOC team members explain the network to attendees in front of each rack.

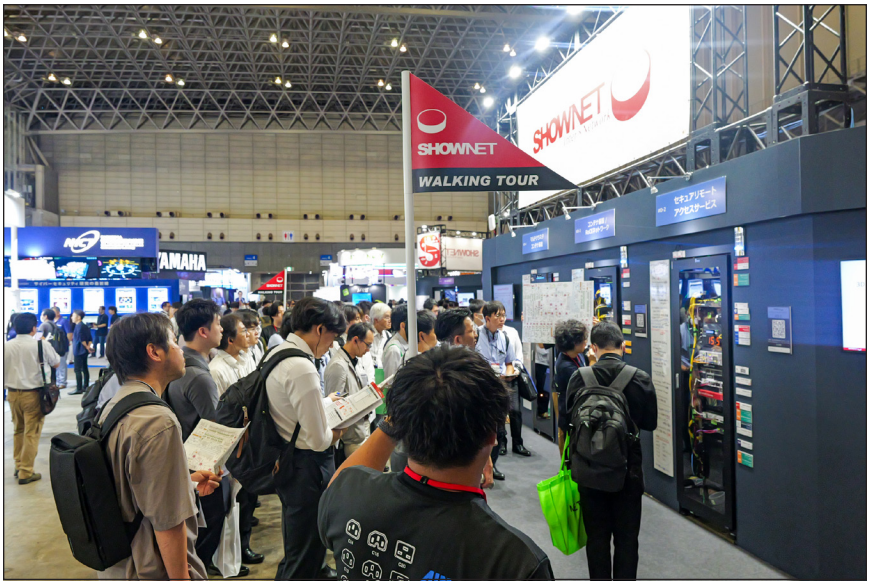

Figure 9: A scene of the tear-down process in 2024. All devices were unmounted from the racks.

References and Further Reading

[0] David Strom, "The Interop Shownet," *The Internet Protocol Journal*, Volume 27, No. 3, October 2024.

[1] Clarence Filsfils, Pablo Camarillo, John Leddy, Daniel Voyer, Satoru Matsushima, and Zhenbin Li, "Segment Routing over IPv6 (SRv6) Network Programming," RFC 8986, February 2021.

[2] Francois Clad, Xiaohu Xu, Clarence Filsfils, Daniel Bernier, Cheng Li, Bruno Decraene, Shaowen Ma, Chaitanya Yadlapalli, Wim Henderickx, and Stefano Salsano, "Service Programming with Segment Routing," Internet-Draft, Work in Progress, February 2025.
draft-ietf-spring-sr-service-programming-09

[3] Ryo Nakamura, Yukito Ueno, and Teppei Kamata, "An experiment of SRv6 Service Chaining at Interop Tokyo 2019 Shownet," Internet-Draft, Work in Progress, October 30, 2019.
draft-upa-srv6-service-chaining-exp-00

[4] Glenn Evans, "Inside InteropNet's Hot Stage,"*Network Computing*, April 2013.

[5] *deadman*: **https://github.com/upa/deadman**

[6] Interop 2024 ShowNet map:
https://www.interop.jp/2024/assets/file/e-web.pdf

[7] ShowNet map icons:
https://github.com/interop-tokyo-shownet/shownet-icons

TAKASHI TOMINE received a Master's degree from Keio University, Japan, and finished his Ph.D. program without a dissertation at Keio University. He is now an Associate Senior Research Engineer at the National Astronomical Observatory of Japan. He has been an Interop Tokyo ShowNet NOC team generalist since 2013. His research interests include network operation, international research educational networks, and cybersecurity. He can be reached at: **tomine@interop-tokyo.net**

RYO NAKAMURA received his Ph.D. degree in Information Science and Technology from the University of Tokyo, Tokyo, Japan, in 2017. He is currently an Associate Professor at the Information Technology Center, the University of Tokyo, where he operates the university's campus network. His research interests include networking in operating systems, network virtualization, and network operations. Since 2009, he has been involved in Interop Tokyo ShowNet, as a ShowNet team member until 2011, and as a member of the NOC team from 2012 to the present. He has been primarily responsible for the backbone network of ShowNet, and led demonstrations of SDN-related technologies from 2013 to 2017. He can be reached at: **ryo@interop-tokyo.net**

RYOTA "ROY" MOTOBAYASHI holds a Bachelor of Engineering from Shinshu University and is qualified as CISSP, Japan's Registered Information Security Specialist and Information Technology Engineer (Class I and Network Specialist). He went through various networking-related projects, from hardware design to corporate strategy planning for NEC Corporation 1988-2023. Since 2024, he has worked for Telecom Engineering Center, a certification body in Japan. His long-term contributions to Interop Tokyo are NOC 1994–1996, NOC Advisory 2006–2017, and Program Committee 2004–2023. He can be reached at: **jj1wtl@jarl.com**

The IPv6 Transition

by Geoff Huston

The state of the transition to IPv6 within the public Internet continues to confound us. RFC 2460[1], the first complete effort at a specification of the IPv6 protocol, was published in December 1998, more than 25 years ago. The entire point of IPv6 was to specify a successor protocol to IPv4 because of the prospect of running out of IPv4 addresses. Yet while the public Internet ran out of IPv4 addresses more than a decade ago, the contrary observation is that this network platform is still largely sustained through its use of IPv4. The transition of the public Internet to the IPv6 protocol has been going on for 25 years now, and if there were any urgency to be instilled in the transition effort by the prospect, and then the reality, of IPv4 address exhaustion, then we've been living with exhaustion a very long time now, and we're largely inured to it. It's probably time to ask the question again: How much longer will this transition to IPv6 take?

At APNIC Labs[0] we've been measuring the uptake of IPv6 for more than a decade now. We use a measurement approach that looks at the network from the perspective of the Internet user base. What we measure is the proportion of users who can reach a published service when the only means to do so is by using IPv6. The data is gathered using a measurement script embedded in an online ad, and the ad placements are configured to sample a diverse collection of end users continually.

Figure 1 displays the IPv6 adoption report showing our measurements of IPv6 adoption across the Internet user base from 2014 to 2024.

Figure 1: IPv6 Adoption – 2014 to 2024. (APNIC Labs Data)

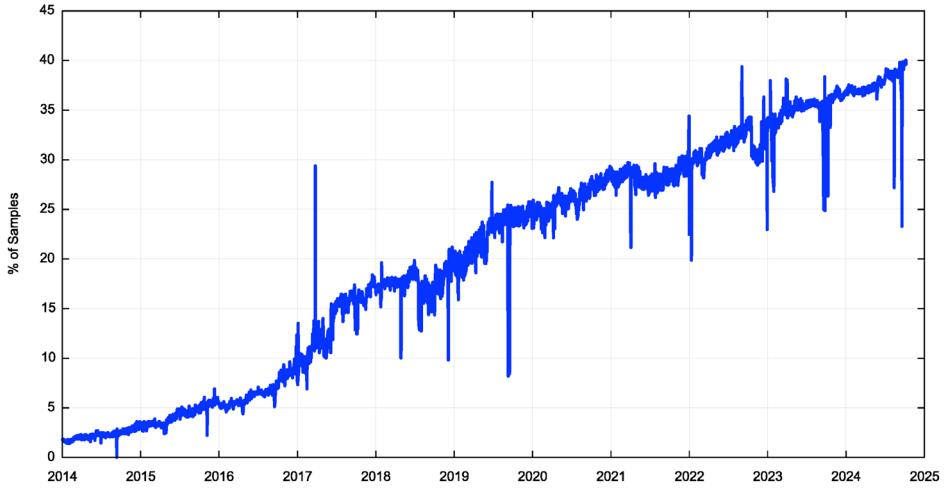

On the one hand, Figure 1 is one of those classic "up and to the right" Internet curves that shows continues growth in the adoption of IPv6. The problem is in the values in the scale of the Y-axis. The issue here is that in 2024 we were at a level where only a little more than one-third of the Internet user base could access an IPv6-only service. Everyone else, now in 2025, is still in an IPv4-only Internet.

This situation appears to be completely anomalous. It's been more than a decade since the supply of "new" IPv4 addresses was exhausted, and the Internet has not only been running on empty, but also is now tasked to span an ever-increasing collection of connected devices—and it has achieved this feat without collapsing. In late 2024 it is variously estimated (or guessed!) that some 20 billion devices used the Internet, yet the Internet IPv4 routing table encompasses only some 3.03 billion unique IPv4 addresses. The original "end-to-end" architecture of the Internet assumed that every device was uniquely addressed with its own IP address, yet the Internet is now sharing each individual IPv4 address across an average of 6 devices, and apparently it all seems to be working! If "end-to-end" was the sustaining principle of the Internet architecture in the 1980's, then as far as the current users of IPv4-based access and services across the public Internet are concerned, it's all over!

IPv6 was meant to address these issues, and the 128-bit wide address fields in the protocol have sufficient address space to allow every connected device to use its own unique address. The design of IPv6 was intentionally very conservative. To a first level of approximation IPv6 is simply "IPv4 with bigger addresses." There are also some changes to fragmentation controls, the address acquisition protocols [*Address Resolution Protocol* (ARP) vs. *Neighbour Discovery*], and the *IP Options* fields, but the upper-level transport protocols are unchanged. IPv6 was intended to be a largely invisible change to a single level in the protocol stack, and definitely not intended to be a massive shift to an entirely novel networking paradigm.

In the sense of representing a very modest incremental change to IPv4, the IPv6 design achieved its objective, but in so doing it necessarily provided little in the way of any marginal improvement in protocol use and performance. IPv6 was no faster, no more versatile, no more secure than IPv4. The major benefit of IPv6 was to mitigate the future risk of IPv4 address exhaustion. In terms of conventional market operations, many markets, including that of the Internet, apply a hefty discount factor to future risk. The result is that the level of motivation to undertake this transition is highly variable given that the expenditure to deploy this second protocol does not immediately realize tangible benefits in terms of lower cost, greater revenue, or greater market share. In a networking context where market-based coordination of individual actions is essential, a level of diversity of views of the net value of running a dual-stack network often leads to reluctance on the part of individual actors and sluggish progress of the common outcome of the transition. As a result, there is no common sense of urgency.

To illustrate this fact, we can look at the time series shown in Figure 1 and ask the question: "If the growth trend of IPv6 adoption continues at its current rate, how long will it take for every device to be IPv6-capable?"

Asking this question is the same as looking at a linear trend line placed over the data series used in Figure 1 for the date when this trend line reaches 100%. Using a least-squares best fit for this data set from January 2020 to the present day, and using a linear trend line, we can come up with Figure 2.

This exercise predicts that we'll see completion of this transition in late 2045, or some 20 years into the future. It must be noted that there is no deep modelling of the actions of various service providers, consumers, and network entities behind this prediction. The only assumption that drives this prediction is that the forces that shaped the immediate recent past are unaltered when looking into the future. In other words, this exercise simply assumes that "tomorrow is going to be a lot like today."

Figure 2: IPv6 Adoption – Projection.

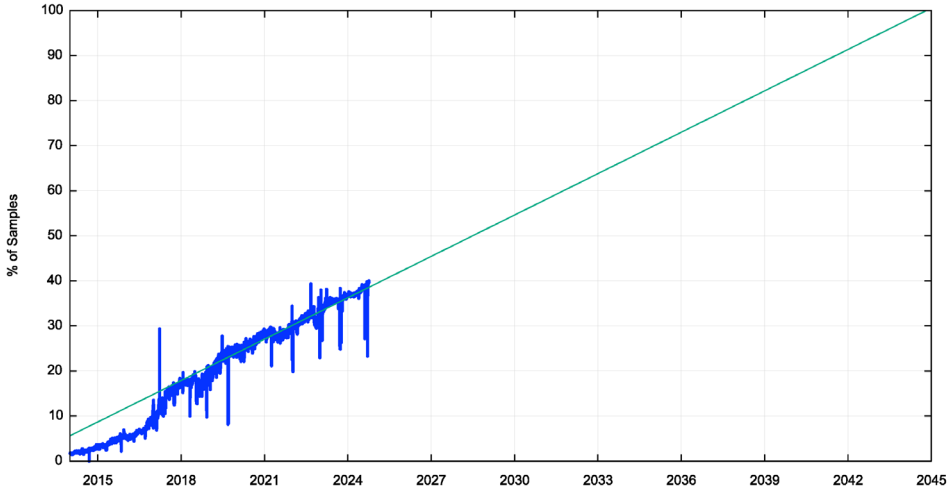

The projected date in Figure 2 is less of a concern than the observation that this model predicts a continuation of this transition for a further two decades. If the entire concept of IPv6 was to restore a coherent address plan across the collection of Internet-connected devices, then placing this model of coherent unique device addressing in abeyance for some 30 years, from around 2015 through to 2045, leads to questioning the role and value of such a unique device addressing framework in the first place! If we can operate a fully functional Internet without such a coherent end-device address architecture for three decades, why would we feel the need to restore address coherence at some point in the future? What's the point of IPv6 if it's not address coherence?

Something has gone very wrong with this IPv6 transition, and that's what I will examine in this article.

A Little Bit of History

By 1990 it was clear that IP had a problem. The Internet was still tiny at the time, but the growth patterns were exponential, doubling in size every 12 months. We were stressing out the pool of Class B IPv4 addresses, and in the absence of any corrective measures this address pool would be fully depleted in 1994 (Figure 3).

Frank Solensky presented predictions at the 18th meeting of the *Internet Engineering Task Force* (IETF)[10].

Figure 3: IPv4 Depletion Predictions, Frank Solensky.

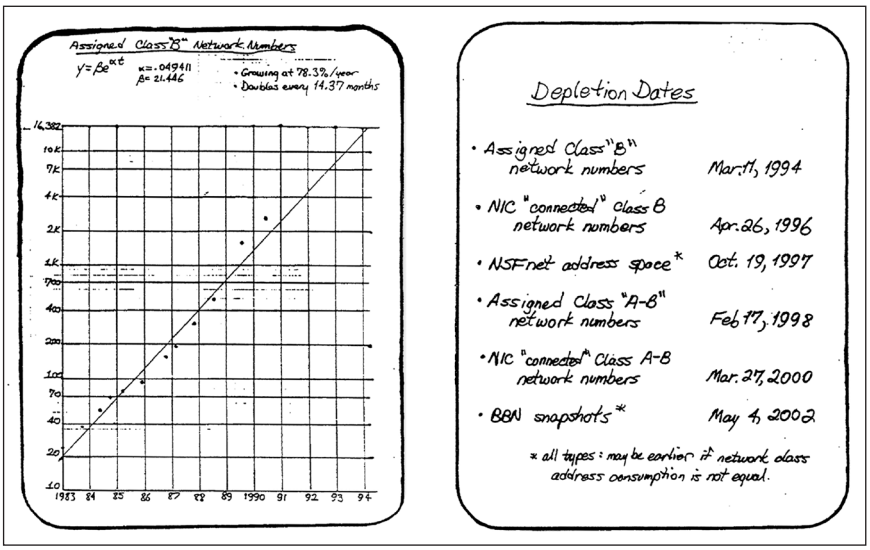

We were also placing pressure on the routing system at the time. The deployed routers in 1992 had only enough memory to support a further 12 to 18 months of routing growth. The combination of these routing and addressing pressures was collectively addressed in the IETF at the time under the umbrella of the ROAD effort, as described in RFC 1380[2].

There was a collection of short-, medium- and longer-term responses that were adopted in the IETF to address the problem. In the short term, the IETF dispensed with the class-based IPv4 address plan and instead adopted a variably sized address prefix model. Routing protocols, including the *Border Gateway Protocol* (BGP), were quickly modified to support these classless address prefixes. Variably sized address prefixes added additional burdens to the address-allocation process, and in the medium term the Internet community adopted the organisational measure of the *Regional Internet Registry* (RIR) structure to allow each region to resource the increasingly detailed operation of address-allocation and registry functions for their region. These measures increased the specificity of address allocations and provided the allocation process with a more exact alignment to determine adequate resource allocations that permitted a more diligent application of relatively conservative address-allocation practices. These measures realized a significant increase in address usage efficiency. The concept of "address sharing" using *Network Address Translation* (NAT)[3] also gained some traction in the *Internet Service Provider* (ISP) world. Not only did NATs dramatically simplify the address administration processes in ISPs, they also played a major role in reducing the pressures on overall address consumption.

The adoption of these measures across the early 1990's pushed a 2-year imminent crisis into a more manageable decade-long scenario of depletion. However, they were not considered to be a stable long-term response. It was thought at the time that an effective long-term response really needed to extend the 32-bit address field used in IPv4. Then the transition from mainframe to laptop was well underway in the computing work, and the prospect of further reductions in size and expansion of deployment in smaller embedded devices was clear. An address space of 4 billion was just not large enough for what was likely to occur in the coming years in the computing world.

But in looking at a new network protocol with a vastly increased address space, the IETF realized that any such change would not be backward-compatible with the installed base of IPv4 systems. As a result, there were a few divergent schools of thought as to what to do. One approach was to jump streams and switch over to use the Connectionless Transport profile of the *Open Systems Interconnection* (OSI) Protocol Suite and adopt OSI *Network Service Access Point Address* (NSAP) addresses along the way. Another was to change as little as possible in IP except the size of the address fields. And numerous ideas were thrown about in the area of proposing significant changes to the IP model.

By 1994 the IETF had managed to settle on the minimal change approach, which was IPv6. The address field was expanded to 128 bits, a *Flow ID* field was introduced, fragmentation behaviour was altered and pushed into an optional header, and ARP was replaced with *multicast*.

The main thing to note was that IPv6 did not offer any new functionality that was not already present in IPv4. It did not introduce any significant changes to the operation of IP. It was just IP with larger addresses.

Transition

While the design of IPv6 consumed a lot of attention at the time, the concept of transition of the network from IPv4 to IPv6 did not.

Given the runaway adoption of IPv4, there was a naive expectation at that time that IPv6 would similarly just take off, and there was no need to give the transition much thought. In the first phase, we would expect to see applications, hosts, and networks adding support for IPv6 in addition to IPv4, transforming the Internet into a dual-stack environment. In the second phase we could then phase out support for IPv4. The expectation was that the process would take a few years.

This plan had numerous problems. Perhaps the most serious one was a resource-allocation problem. The Internet was growing extremely quickly, and most of our effort was devoted to keeping pace with demand. More users, more capacity, larger servers, more content and services, more responsive services, more security, better defence. All of these factors shared a common theme: *scale*.

We could either concentrate our resources on meeting the incessant demands of scaling, or we could work on IPv6 deployment. The short- and medium-term measures that we had already taken had addressed the immediacy of the problems of address depletion, so in terms of priority, scaling was a for more important priority for the industry than IPv6 transition. Through the decade from 1995 to 2005 the case for IPv6 quietly slumbered in terms of mainstream industry attention.

IPv4 addresses were still available, and the use of *Classless Inter-Domain Routing* (CIDR) and far more conservative address-allocation practices had pushed the prospect of IPv4 address depletion out by more than a couple of decades. Many more pressing operational and policy issues for the Internet absorbed the industry's collective attention in those days.

However, this period of respite was brief. The scaling problem accelerated by a whole new order of magnitude in the mid 2000's with the introduction of the iPhone and its brethren[4]. Suddenly this scale problem was not just of the order of tens or even hundreds of millions of households and enterprises, it transformed into a problem of billions of individuals and their personal devices, and it added mobility into the mix. As a taste of a near-term future, the production scale of these "smart" devices quickly ramped up into annual volumes of hundreds of millions of units. The entire reason why IPv6 was a necessity was coming into fruition, but at this stage we were just not ready to deploy IPv6 in response. Instead, we rapidly increased our consumption of the remaining pools of IPv4 addresses and we supported the first wave of large-scale mobile services with IPv4. Dual stack was not even an option in the mobile world at the time. The rather bizarre economics of financing 3G infrastructure meant that dual-stack infrastructure in a 3G platform was impractical, so IPv4 was used to support the first wave of mobile services. This situation quickly turned to IPv4 and NATs as the uptake of mobile services gathered momentum.

At the same time the decentralised nature of the Internet was hampering IPv6 transition efforts. What point was there in developing application support for IPv6 services if no host had integrated IPv6 into its network stack? What point was there in adding IPv6 to a host networking stack if no ISP was providing IPv6 support? And what point was there for an ISP to deploy IPv6 if no hosts and no applications would use it? In terms of IPv6 at this time, nothing happened.

The operating-system sector made the first efforts to try to break this impasse of mutual dependence, and fully functional IPv6 stacks were added to the various flavours of Linux, Windows, and MacOS, as well as in the mobile host stacks of iOS and Android.

But even these implementations were not enough to allow a transition to achieve critical momentum. It could be argued that this situation made the IPv6 situation worse and set back the transition by some years.

The problem was that with IPv6-enabled hosts there was some desire to use IPv6, but these hosts were isolated "islands" of IPv6 sitting in an ocean of IPv4. The concentration of the transition effort then fixated on various tunnelling methods to tunnel IPv6 packets through the IPv4 networks (Figure 4). While you can perform this tunnelling manually when you have control over both tunnel endpoints, this approach was not that useful. What we wanted was an automated tunnelling mechanism that took care of all these details.

Figure 4: Phase 1 of the IPv6 Transition.

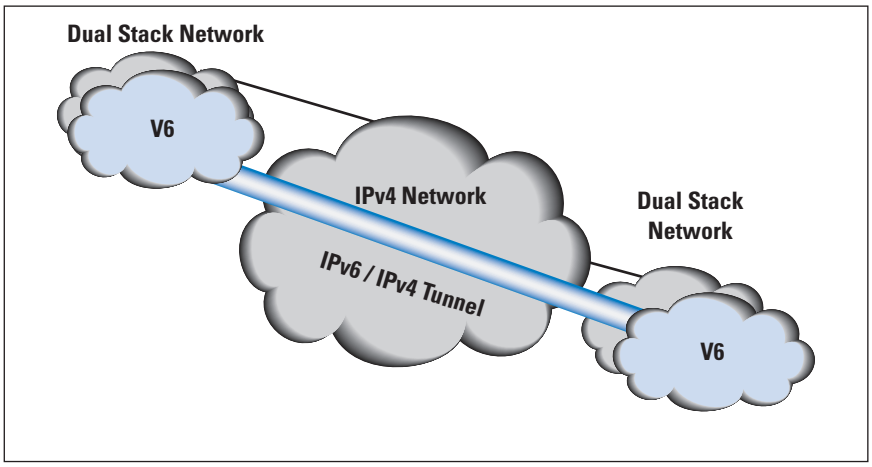

The first such approach that gathered some momentum was 6to4[5]. The first problem with 6to4 was that is required public IPv4 addresses, so it could not provide services to IPv6 hosts that were behind a NAT. The more critical problem was that firewalls had no idea how to handle these 6to4 packets, and the default action when in doubt is to deny access. So 6to4 connections encountered an average of a 20 to 30% failure rate in the public Internet, making it all but unusable as a mainstream service. The NAT traversal issue was also a problem, so a second auto-tunnel mechanism was devised that performed NAT sensing and traversal. This mechanism, *Teredo*[6], was even worse in terms of failure rates, and some 40% of Teredo connection attempts were observed to fail[7].

Not only were these Phase 1 IPv6 transition tools extremely poor performers, as they were so unreliable, but even when they worked the connection was both fragile and slower than IPv4. The result was perhaps predictable, even if unfair. It was not just the transition mechanisms that were viewed with disfavour, but IPv6 itself also attracted some opprobrium.

Up until around 2011 IPv6 was largely ignored as a result in the mainstream of the public Internet. A small number of service providers tried to deploy IPv6, but in each case they found themselves with a unique set of challenges that they and their vendors had to solve, and without a rich set of content and services on IPv6, the value of the entire exercise was highly dubious! So, nothing much happened.

Movement at Last!

It wasn't until the central IPv4 address pool that the *Internet Assigned Numbers Authority* (IANA) managed was depleted at the start of 2011, and the first RIR, APNIC, ran down on its general allocation pool in April of that year, that the ISP industry started to pay more focussed attention to this IPv6 transition.

At around the same time, the mobile industry commenced its transition into 4G services. The essential difference between 3G and 4G was the removal of the *Point-to-Point Protocol* (PPP) tunnel through the radio access network from the gateway to the device and its replacement by an IP environment. This solution allowed a 4G mobile operator to support a dual-stack environment without an additional cost component, and it was a major enabler for IPv6. Mapping IPv4 into IPv6 (or the reverse) is fragile and inefficient for service providers as compared to native dual stack. In the 6-year period from 2012 to the start of 2018, the level of IPv6 deployment rose from 0.5 to 17.4%. At this stage IPv6 was no longer predominately tunnelled, as many networks supported IPv6 in native node (Figure 5).

Figure 5: Phase 2 of the IPv6 Transition.

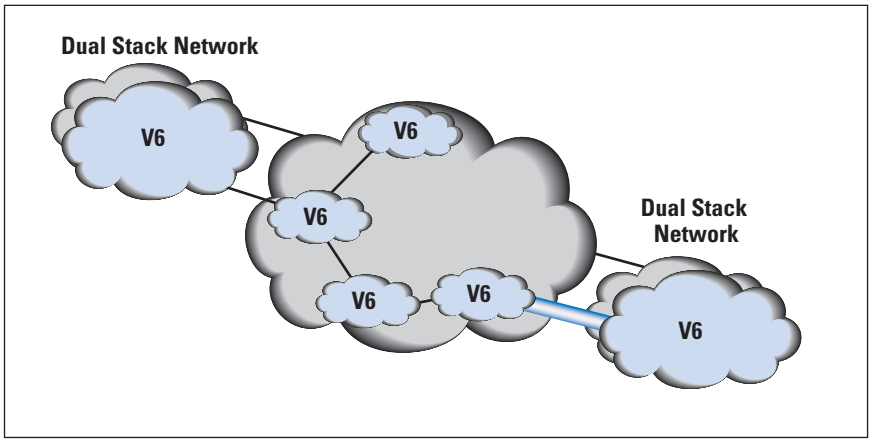

The problem here was that we were late with this phase of the transition. The intention of this transition was to complete the work and equip every network and host with IPv6 before we ran out of IPv4 addresses (Figure 6).

The position we had arrived at by 2012 was far more challenging. The pools of available IPv4 address space were rapidly depleting, and the regional address policy communities were introducing highly conservative address-allocation practices to eke out the remaining address pools. At the same time the amount of IPv6 uptake was minimal. The transition plan for IPv6 was largely broken (Figure 7).

Figure 6: The IPv6 Transition Plan.

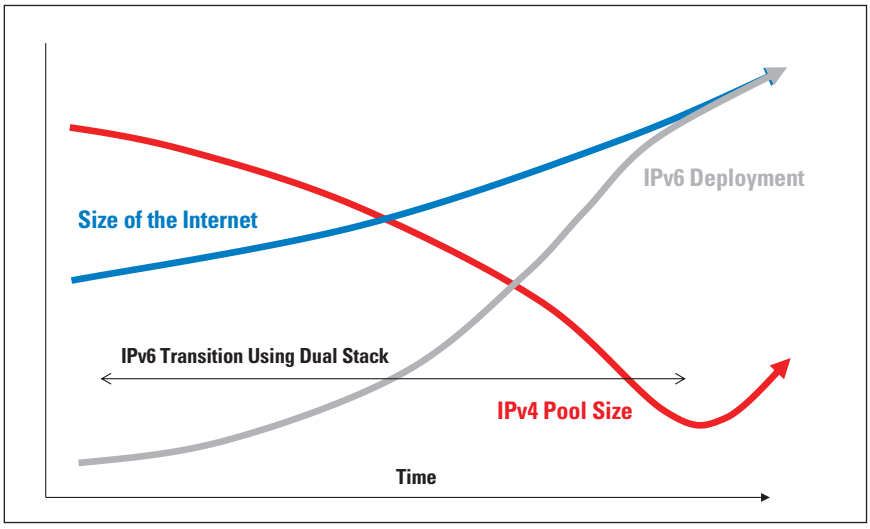

Figure 7: The IPv6 Transition Plan in 2012.

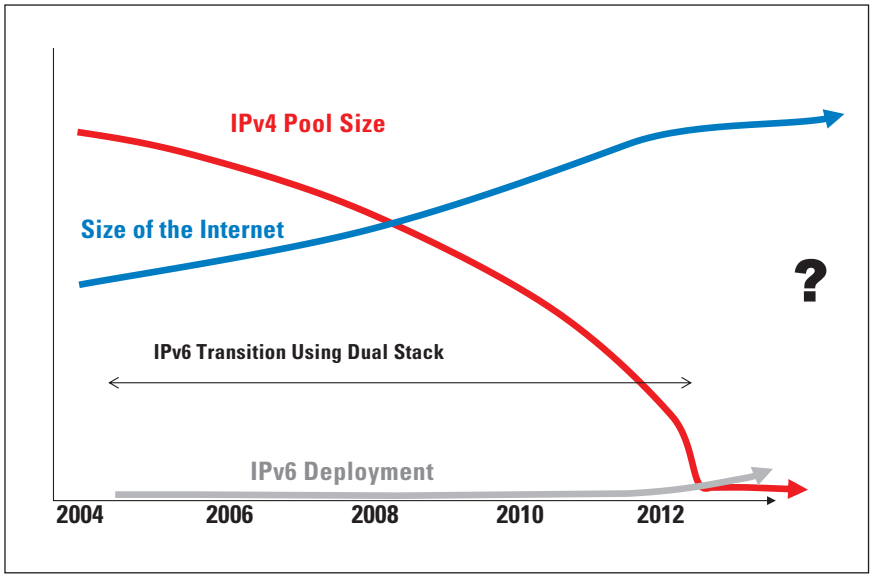

NATs and Address Scarcity Pressures

At this point there was no choice for the Internet, and to sustain growth in the IPv4 network while we were waiting for IPv6 to gather momentum we turned to NATs. NATs were a challenging subject for the IETF. The entire concept of coherent end-to-end communications was to eschew active middleware such as NATs in the network. NATs created a point of disruption in this model, thereby causing a critical dependency upon network elements. They removed elements of network flexibility from the network and at the same time reduced the set of transport options to the *Transmission Control Protocol* (TCP) and *User Datagram Protocol* (UDP).

The IETF resisted any efforts to standardise the behaviour of NATs, fearing perhaps that standard specifications of NAT behaviour would bestow a legitimacy on the use of NATs, an outcome that many IETF participants were very keen to avoid.

This aversion did not reduce the level of impetus behind NAT deployment. We had run out of IPv4 addresses and IPv6 was still a distant prospect, so NATs were the most convenient solution. What this action did achieve was to create a large variance of NAT behaviours[15] in various implementations, particularly with respect to UDP behaviours. This situation has exacted a cost in software complexity where an application needs to dynamically discover the type of NAT (or NATs) in the network path if it wants to perform anything more complex than a simple two-party TCP connection.

Despite these issues NATs were a low-friction response to IPv4 address depletion, where individual deployment could be undertaken without incurring external dependencies. On the other hand, deployment of IPv6 was dependant on other networks and servers also deploying IPv6. NATs made highly efficient use of address space for clients, as not only could a NAT use the 16-bit source port field, but by time-sharing the NAT binding, NATs achieved an even greater level of address efficiency. A major reason why we've been able to sustain an Internet with tens of billions of connected devices is through the widespread use of NATs.

Server architectures were also changing. The introduction of *Transport Layer Security* (TLS)[8] into the web-server world included a point in TLS session establishment where the client informs the server platform the name of the service that it intends to connect to. Not only did this information allow TLS to validate the authenticity of the service point, but it also allowed a server platform to host an extremely large collection of services from a single platform (and a single platform IP address) and perform individual service selection via this TLS *Server Name Indication* (SNI). The result is that server platforms perform service selection by name-based distinguishers [*Domain Name System* (DNS) names] in the session handshake, allowing a single server platform to serve large numbers of individual servers. The implications of the widespread use of NATs and the use of server sharing in service platforms has taken the pressure off the entire IPv4 address environment.

One of the best ways to illustrate the changing picture of address scarcity pressure in IPv4 is to look at the market price of address transfers over the past decade. Scarcity pressure is reflected in the market price. Figure 8 shows a time series of the price of traded IPv4 addresses.

The period of the COVID outbreak coincided with a rapid price escalation over 2021, but the price has since declined to between $30 and $40 per address, and this price, admittedly over a $16 range from $26 to $42 per address, was stable across 2024. This price data indicates that IPv4 addresses were still in demand in 2024, but the level of demand appears to have equilibrated against available levels of supply, implying that there was no scarcity premium in evidence in the address market in 2024. This data points to the combination of the efficacy of NATs in extending the efficiency of IPv4 addresses by using the 16 bits of port address space plus the additional benefits of using shared address pools.

Figure 8: Market Prices of IPv4 Address Transfers. (Data from Hilco Streambank)

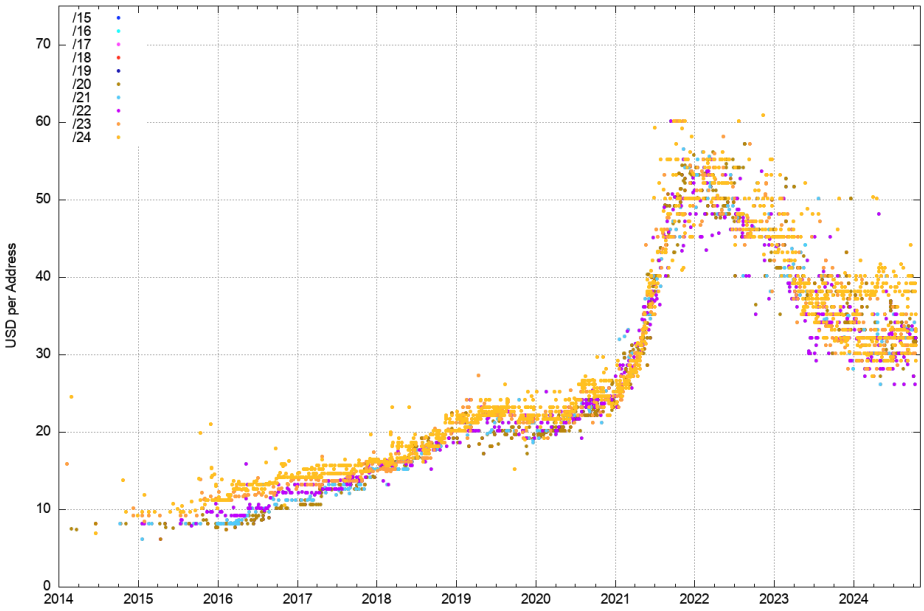

However, it's not just IPv4 that has alleviated the scarcity pressure for IPv4 addresses. Figure 1 indicates that over the past decade the level of IPv6 adoption has risen to encompass some 40% of the user base of the Internet. Most applications, including browsers, support *Happy Eyeballs*[9], which is a shorthand notation for preferring to use IPv6 over IPv4 if both protocols are available for use in support of a service transaction. As network providers roll out IPv6 support, the pressure on their IPv4 address pools for NAT use is relieved because the applications prefer to use IPv6 where available.

How Much Longer?

Now that we are somewhere in the middle of this transition, the question is: How much longer is this transition going to take?

This question seems simple, but it does need a little more elucidation. What is the "end point" when we can declare the transition to be over? When will this transition be "complete"? Is it the time when the Internet has no more IPv4-based traffic? Or is it the time when the Internet no longer requires IPv4 in public services? Or do we mean the point when IPv6-only services are viable? Or perhaps we should look at the market for IPv4 addresses and define the endpoint of this transition at the time when the price of IPv4 addresses completely collapses? Perhaps we can take a more pragmatic position here and rather than looking for completion as the point when the Internet is completely bereft of all use of IPv4 addresses and their use, we could define "completion" as the point when use of IPv4 is no longer necessary. The implication would be that when a service provider can operate a viable Internet service using only IPv6 and having no supported IPv4 access mechanisms at all, then we would have completed this transition.

What is the implication? Certainly, the ISP needs to provide IPv6. But all the connected edge networks and the hosts in these networks need to support IPv6 as well. After all, the ISP has no IPv4 services at this point of completion of the transition. It also implies that all the services the clients of this ISP use must be accessible over IPv6. Yes, this accessibility includes all the popular cloud services and cloud platforms, all the content streamers, and all the content-distribution platforms. It also includes specialised platforms such as Slack, Xero, Atlassian, and similar platforms. The data published at the Internet Society's *Pulse* page[11] reports that only some 47% of the top 1000 web sites are reachable over IPv6, so clearly a lot of service platforms have work to do, and this work will take more time.

When we look at the IPv6 adoption data for the United States, another somewhat curious anomaly is evident (Figure 9).

Figure 9: IPv6 Adoption in the US - 2014 to 2024. (APNIC Labs Data)

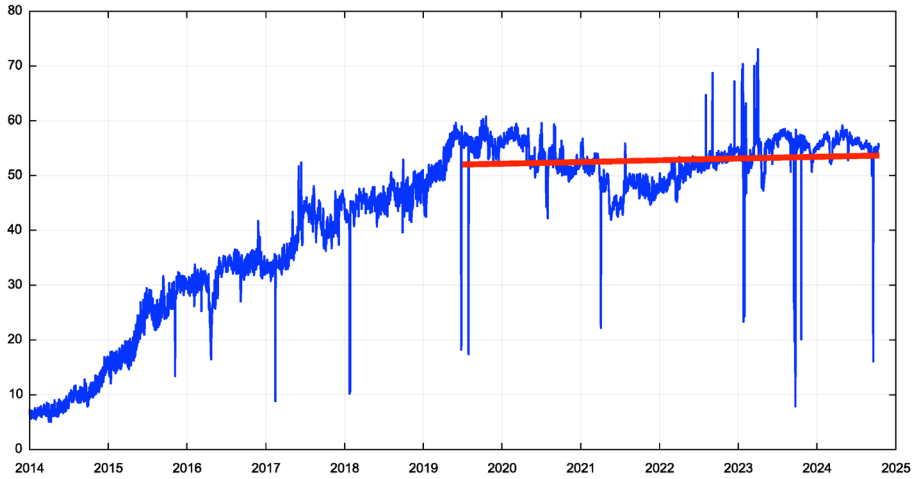

The data shows that the level of IPv6 use in the US has remained constant since mid-2019. Why is there no further momentum to continue with the transition to IPv6 in this part of the Internet? I would offer the explanation that the root cause is a fundamental change in the architecture of the Internet.

Changes to the Internet Architecture

The major change to the Internet architecture is a shift away from a strict address-based architecture. Clients no longer need to use a persistent unique public IP address to communicate with servers and services. And servers no longer need to use a persistent unique public IP address to provide clients with access to the service or content. Address scarcity takes on an entirely different dimension when unique public addresses are not required to number every client and every distinct service.

Some of the clues that show the implications of this architectural shift are evident when you look at the changes in the internal economy of the Internet. The original model of IP was a network protocol that allowed attached devices to communicate with each other.

The network providers supplied the critical resource to allow clients to consume content and access services. At the time the costs of the network service dominated the entire cost of the operation of the Internet, and in the network domain distance was the dominant cost factor.

Network providers who offered distance services (so-called "transit providers") were the dominant ones. Little wonder that we spent a lot of our time working through the issues of interconnection of network service providers, customer/provider relationships, and various forms of peering and exchanges. The ISPs were in effect brokers in the rationing of the scarce resource of distance capacity. This economy was a classic network economy (Figure 10).

Figure 10: The Classic Network Economy.

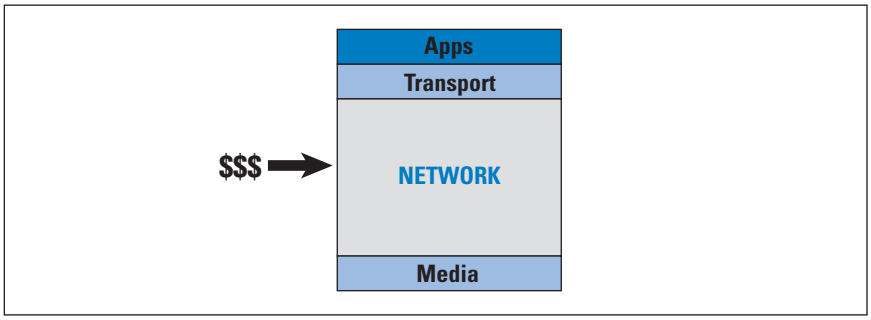

For many years the demand for communications services outstripped available inventory, and price was used as a distribution function to moderate demand against available capacity. However, everything changed because of the effects of *Moore's Law* consistently changing the cost of computing and communications.

The most obvious change has been in the count of transistors in a single integrated circuit. Figure 11 shows the transistor count over time since 1970.

The latest production chips at the end of 2024 were the Apple M3, a 3nm chip with up to 92 billion transistors. With perhaps the possible exception of powering AI infrastructure, these days processing capability is an abundant and cheap resource.

This continual refinement of integrated-circuit production techniques affects the size and unit cost of storage (Figure 12). While the speed of memory has been relatively constant for more than a decade, the unit cost of storage has been dropping exponentially for many decades. Storage is also an abundant resource.

These changes in the capabilities of processing have also profoundly affected communications costs and capacities. The constraining factor in fibre communications systems is the capabilities of the digital signal processors and the modulators. As silicon capabilities improve, it's possible to improve the signal-processing capabilities of transmitters and receivers, which allows for a greater capacity per wavelength on a fibre circuit (Figure 13).

Figure 11: Transistor Count over Time[12].

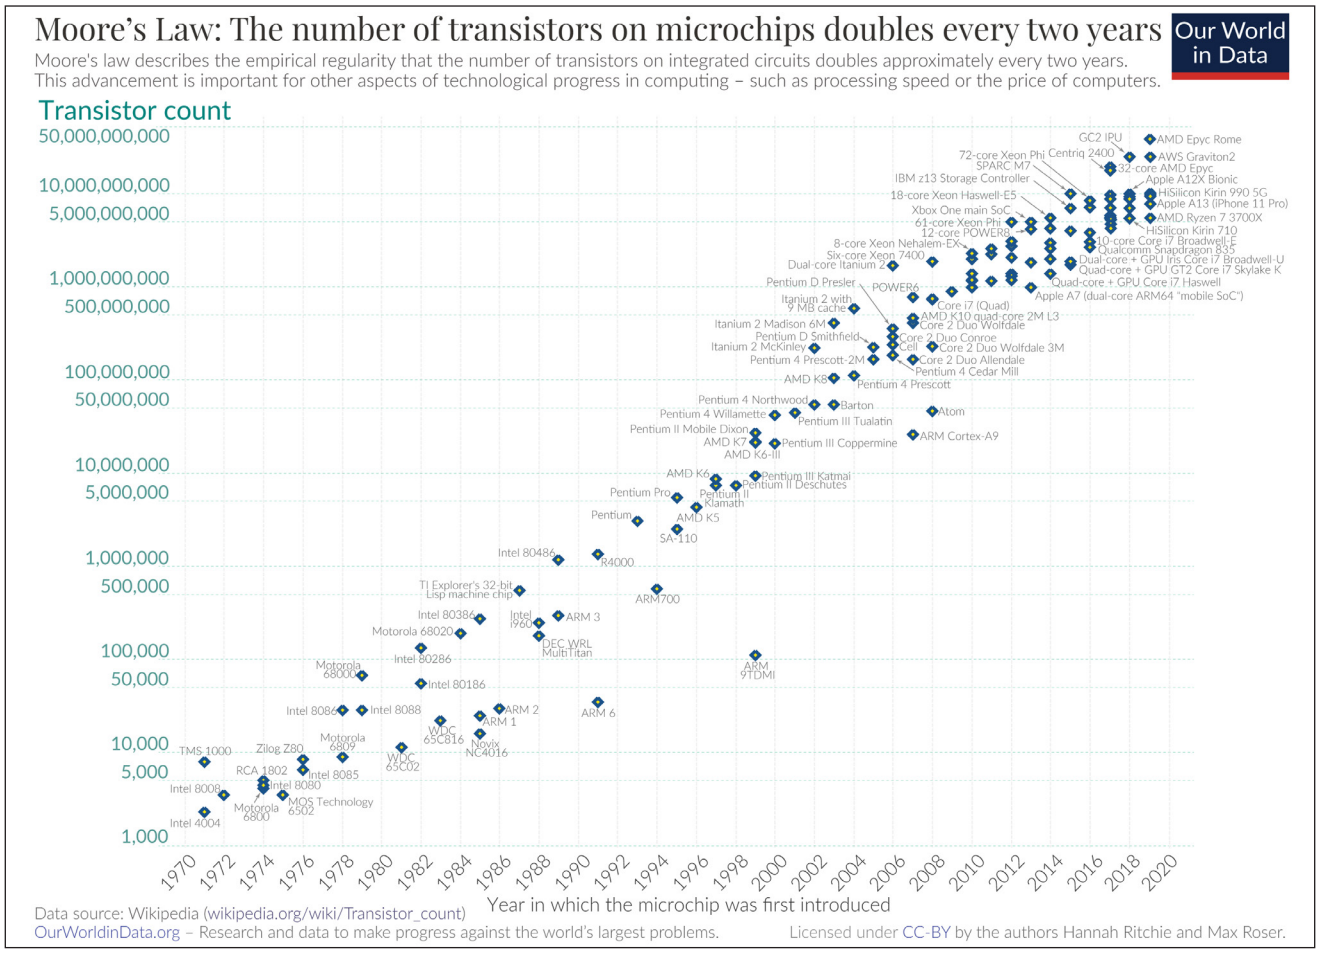

This change from scarcity to abundance in processing, storage, and transmission capacity has had a profound impact on the service model of the Internet. The model has changed from an *on-demand pull* to a *just-in-case* model of pre-provisioning. These days we load replicas of content and services close to the edge of the network where the users are located and attempt to deliver as much of the content and service as possible from these edge points of presence to the users in the adjacent access networks. These changes in the underlying costs of processing and storage have provided the impetus for the expansion of various forms of *Content Distribution Networks* (CDNs), which now serve almost the entirety of Internet content and services. These expansions have allowed us to eliminate the factor of distance from the network, and most network transactions occur over short spans.

The overall result of these changes is the elimination of distance in pushing content and services to clients. We are able to exploit the potential capacity in 5G mobile networks without the inefficiencies of operating the transport protocol over a high-delay connection. Today's access networks operate with greater aggregate capacity, and the close proximity of service-delivery platform and client allow transport protocols to use this capacity, as transport sessions that operate over a low-latency connection are also far more efficient. Service interactions across shorter distances using higher-capacity circuitry results in a much faster Internet!

Figure 12: Computer Memory and Storage Unit Costs over Time[13].

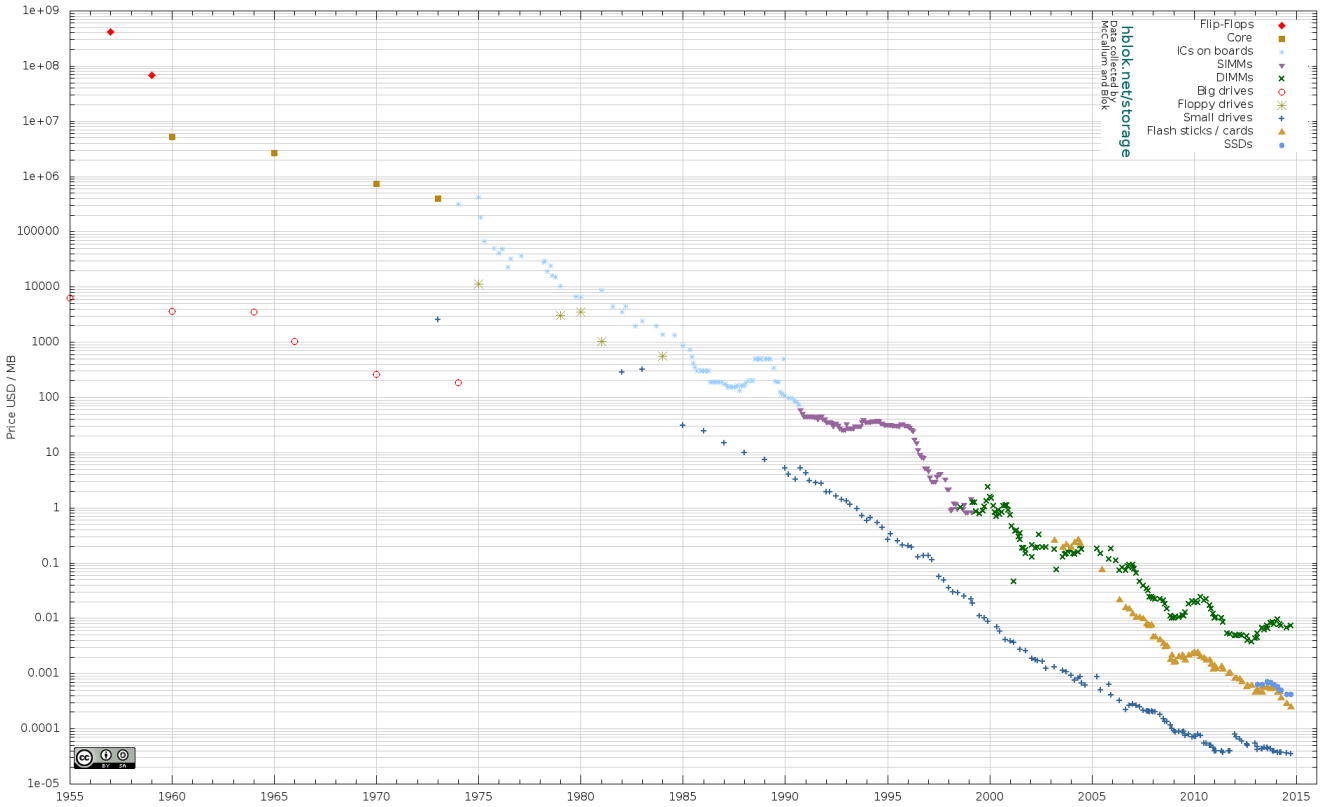

Figure 13: Fibre Capacity over Time[14].

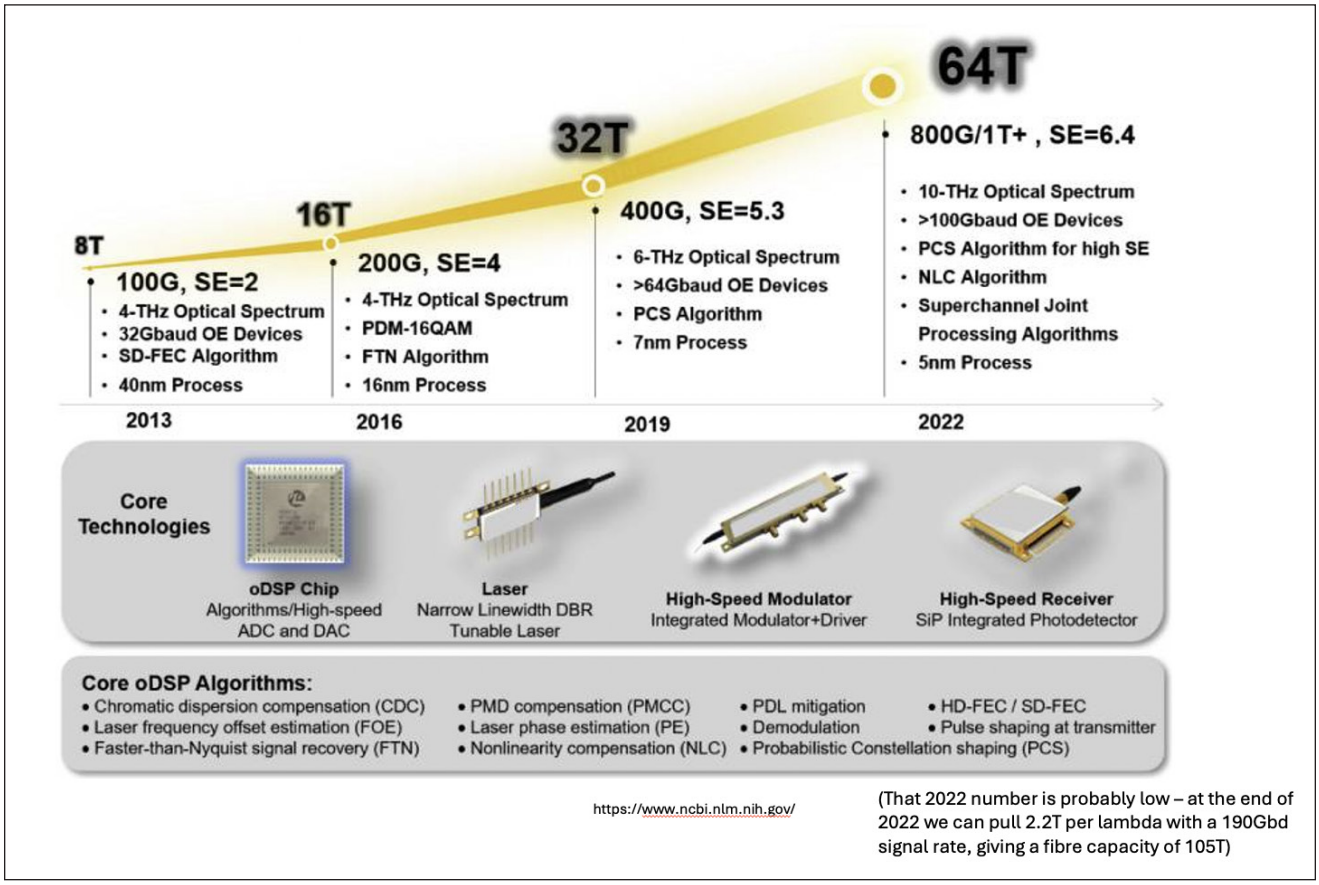

In addition to being bigger and faster, this environment of abundant communications, processing, and storage capacity is operating in an industry that enjoys significant economies of scale. And much of this environment is funded by capitalising a collective asset that is infeasible to capitalise individually, namely the advertisement market. The result of these changes is that a former luxury service accessible to just a few has been transformed into an affordable mass-market commodity service available to all.

However, it's more than just bigger, faster, and cheaper. This shift into abundance of basic inputs for the digital environment has changed the economics of the Internet as well. The role of the network as the arbiter of the scarce resource of communication capability has dissipated. In response, the economic focus of the Internet economy has shifted up the protocol stack to the level of applications and services (Figure 14).

Figure 14: The Transformation of the Network Economy.

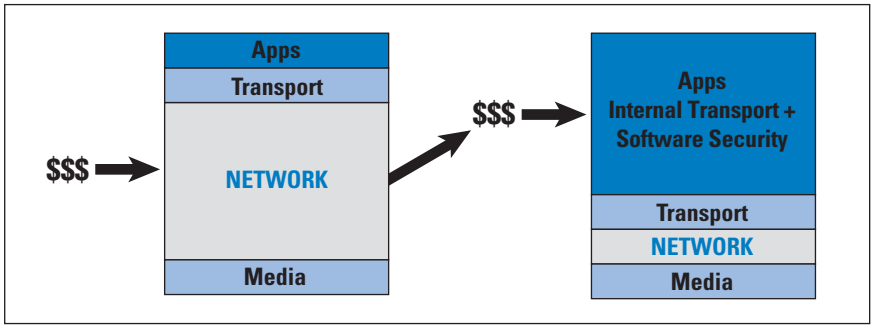

Now let's return to the situation of the transition to IPv6. It is left to networks and network operators to make the investments to switch to a dual-stack platform initially (and then ultimately to remove support for IPv4). But this change is really not visible, or even crucial, to the content or service world. If IPv4 and NATs perform the carriage function adequately, then there is no motivation for the content and service operators to pay a network a premium to have a dual-stack platform.

It's domain names that operate as service identifiers, it's domain names that underpin the users' tests of authenticity of the online service, and it's the DNS that increasingly is used to steer users to the "best" service-delivery point for content or service. From this perspective addresses, IPv4 or IPv6, are not the critical resource for a service and its users. The "currency" of this form of CDN networking is *names*.

So where are we in 2025? Today's public Internet is largely a service-delivery network using CDNs to push content and service as close to the user as possible. The multiplexing of multiple services onto underlying service platforms is an application-level function tied largely to TLS and service selection using the SNI field of the TLS handshake. We use the DNS to perform "closest match" service platform selection. It's the objective of a CDN to directly attach to the access networks where its users are located, and the result is a BGP routing table inside the CDN with an average *Autonymous System* (AS) *Path Length* that is intended to converge to 1!

From this respect the DNS has supplanted the role of routing! We may not route "names" in today's Internet, but it is certainly operating in a way that is largely isomorphic to such a named data network.

This architectural change has a few additional implications for the Internet. TLS, like it or not (and there is much to criticise about the robustness of TLS), is the sole underpinning of authenticity in the Internet. *Domain Name System Security Extensions* (DNSSEC) has not gathered much momentum to date. The protocol is too complex, too fragile, and just too slow to use for most services and their users. Some value its benefits highly enough that they are prepared to live with its shortcomings, but that's not the case for most name holders and most users, and no amount of passionate exhortations about DNSSEC will change this situation! It supports the view that it's not the mapping of a name to an IP address that's critical. What is critical is that the named service can demonstrate that it operated by the owner of the name. Secondly, *Resource Public Key Infrastructure* (RPKI), the framework for securing information being passed in the BGP routing protocol, is really not all that useful in a service network where there is no routing!

The implication of these observations is that the transition to IPv6 is progressing very slowly not because this industry is chronically stupid or short-sighted—something else is going on here. IPv6 alone is not critical to a large set of end-user service-delivery environments. We've been able to take a 1980's address-based architecture and scale it more than a billion-fold by altering the core reliance on distinguisher tokens from addresses to names. There was no real lasting benefit in trying to leap across to just another 1980's address-based architecture (with only a few annoyingly stupid differences, apart from longer addresses!).

Where are we heading in the longer term? We are pushing everything, including value itself, out of the network and over to applications. Transmission infrastructure is becoming an abundant commodity. Network-sharing technology (multiplexing) is decreasingly relevant. We have so much network and computing resources available to us that we no longer have to take consumers to service-delivery points. Instead, we are taking services towards consumers and using the content frameworks to replicate servers and services. With so much computing and storage, the application is becoming the service, rather than just a window to a remotely operated service.

If that's the case, then will networks matter anymore? The last couple of decades have seen us stripping out network-centric functionality and replacing it with an undistinguished commodity packet-transport medium. It's fast and cheap, but it's up to applications to overlay this common basic service with their own requirements. As we push these additional functions out to the edge and ultimately off the network altogether, we are left with simple dumb packet pipes!

You could argue that this situation is nothing new, and it's a continuation of the disruption that the Internet itself brought to bear on the predecessor telephone network infrastructure. The Internet architecture shifted functionality out of the core of the network and replaced synchronous real-time end-to-end virtual circuits with an extremely basic data packet-delivery service where networks were permitted to drop, duplicate, reorder, and re-time these packets in flight across the network.

It was left to the control functions that were embedded in the attached devices (such as the TCP protocol, for example) to create a functional, reliable, end-to-end communications service model. Internet hosts valued a network only to the level of a basic (and imperfect) packet-delivery service. Clients of a network were unwilling to pay a price premium for network-level services that were already being provided by the edge devices.

The result is a diminished network, dramatically reduced in both role and value. This diminished role impairs network operators to raise additional revenue through augmented services, whether it's through variable service responses through *Quality of Service* (QoS) responses or even as basic as IPv6 protocol support.

At this point it's useful to ask: What "defines" the Internet? Is the classic response, namely: "A common shared transmission fabric, a common suite of protocols, and a common protocol address pool" still relevant these days? Or is today's network more like: "A disparate collection of services that share common referential mechanisms using a common name space?"

When we think about what's important to the Internet these days, is the choice of endpoint protocol addressing really important? Is universal unique endpoint addressing a 1980's concept whose time has come and gone? If network transactions are localised, then what is the residual role of unique global endpoint addressing for clients or services? And if we cannot find a role for unique endpoint addressing, then why should we bother? Who decides when to drop this concept? Is this a market function, so that a network that uses local addressing can operate from an even lower cost base to gain a competitive market edge? Or are carriage services so cheap already that the relative benefits in discarding the last vestiges of unique global addresses are so small that it's just not worth bothering about?

And while we ponder such questions, what is the role of referential frameworks in networks? Without a common referential space, how do we usefully communicate? What do we mean by "common" when we think about referential frameworks? How can we join the "fuzzy" human language spaces with the tightly constrained deterministic computer-based symbol spaces?

Certainly, there is much to think about here!

And where does this situation leave the transition to IPv6?

I suspect that the dual-stack world we're in is a world we will be stuck in for quite some time. There seems to be no appetite to resolve this situation by completing the transition any time soon, and absolutely no desire to back out and revert to a IPv4-only network. We are here now, caught in a partial state of transition to IPv6 that is taking on an unfortunate air of permanence! And as the preponderance of value in this environment continues to move up the protocol stack into service, content, and today generative content in the guise of AI, there is little continued capacity to place collective attention on questions that have been left unresolved for decades.

It may well be that the question of when this IPv6 transition will end is a question that engenders decreasing levels of interest and attention in line with the larger picture of the decreasing relative economic value of the answer! Silicon abundance has enabled a few select content and service operators to privatise much of the former public communications platform, and in so doing they have managed to shrink the public Internet to a set of margins at the edges. That reality implies that the answer to the IPv6 transition question may soon be: "Who cares anyway?"

Disclaimer

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

References and Further Reading

[0] APNIC Labs Measurements and Data: **https://labs.apnic.net/measurements/**

[1] Steve Deering and Robert Hinden, "Internet Protocol, Version 6 (IPv6) Specification," RFC 2460, December 1998.

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

[3] Paul Tsuchiya and Tony Eng, "Extending the IP Internet Through Address Reuse," ACM SIGCOMM *Computer Communications Review*, Volume 23, Issue 1, January 1993.

[4] John Laugesen and Yufei Yuan, "What Factors Contributed to the Success of Apple's iPhone?," *Ninth International Conference on Mobile Business and 2010 Ninth Global Mobility Roundtable* (ICMB-GMR), Athens, Greece, 2010, pp. 91–99, DOI: 10.1109/ICMB-GMR.2010.63.

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

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

[7] Geoff Huston, "Stacking it Up," presentation to the IPv6 Operations Working Group, IETF 80, March 2011, **https://www.potaroo.net/presentations/2011-03-31-dualstack.pdf**

[8] Tim Dierks and Christopher Allen, "The TLS Protocol Version 1.0," RFC 2246, January 1999.

[9] Dan Wing and Andrew Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts," RFC 6555, April 2012.

[10] Frank Solensky's Presentation at IETF 18: **https://www.ietf.org/proceedings/18.pdf**

[11] Internet Society *Pulse*: **https://pulse.internetsociety.org/**

[12] Transistor Count over Time: **https://assets.ourworldindata.org/uploads/2020/11/Transistor-Count-over-time.png**

[13] Computer Memory and Storage Unit Costs over Time: **http://aiimpacts.org/wp-content/uploads/2015/07/storage_memory_prices_large-_hblok.net_.png**

[14] Fibre Capacity over Time: **https://www.ncbi.nlm.nih.gov/**

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

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

Book Review

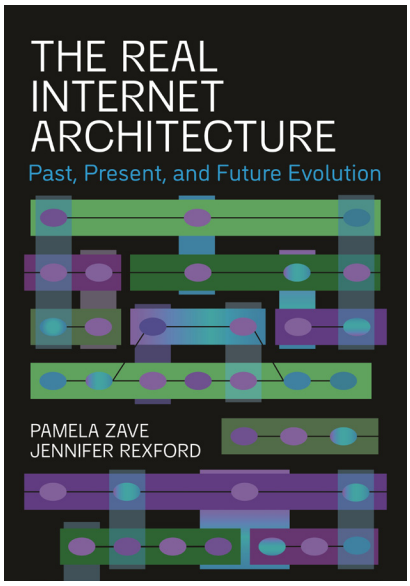

The Real Internet Architecture: Past, Present, and Future Evolution, by Pamela Zave and Jennifer Rexford, Princeton University Press, ISBN 9780691255804, June 2024.

The goal of this book is to present a better way to describe the architecture of various networks, most notably the Internet. The initial portions of the book observe several deficiencies in how we both teach and talk about *networks* today. Most networking courses teach the bits and packets of the Internet, without a unifying framework. Networking papers suffer from "the lack of precise and consistent terminology."

The authors present an alternative architectural model, in the hopes it will help with these issues. The core building block of the model is a network, where a network is defined as having members (hardware and software dedicated to participating in the network), names for members and groups of members, links, topology, a single administrator, and the capacity to route sessions of information. Networks offer users a simple service, namely the ability to send and receive information.

Networks can be connected three different ways:

- *Bridged Networks* are peers. Obviously IEEE 802 Ethernet networks are bridged. In this model, so too are IP networks run by different administrators. The Internet is a bridged (versus routed) network.

- *Layered Networks* put one network on top of another. A *Virtual Private Network* (VPN) is its own network and is placed on top of a set of bridged IP networks. IP networks are layered on bridged link-layer networks. As the examples illustrate, a network may be layered across multiple underlying (bridged) networks.

- Finally, there's the case where a member in a VPN is engaged in a session with a member on a bridged IP network at the layer below the VPN. For that case, the authors repurpose the term *subduction* to describe interactions that cross layers.

This brief description summarizes a much richer conceptual framework in the book, but one can see that a large set of complex network interactions are simplified by putting a box around large chunks, declaring those chunks individual networks (a VPN, an Ethernet, etc.), and then using bridging, layering, and subduction to put them together.

In most cases, the simplification is a relief. The complex interactions among a 5G network, VPNs, and the larger Internet to serve a web-request on my phone become conceptually simpler. Similarly, tenant networks in data centers feel more tractable.

But sometimes the model falters. Treating a *Hypertext Transfer Protocol* (HTTP) session as its own network, as the book does, with a single point-to-point link as its topology, feels simplistic. It is also not at all clear who the single administrator of the HTTP session is (also an issue for some other networks described).

There are also some missed opportunities. I would suggest, in the authors' model, that there are subductive control protocols, of which the *Address Resolution Protocol* (ARP) is probably the most notable, and that these protocols present distinct challenges not encountered by protocols that stay within their network box.

The book is a thought-provoking read. I would be surprised if it managed to persuade many instructors to teach using its paradigms. I think a full textbook would be required to make that happen. At the same time, I think the notion of subduction as a way to reference the challenge of cross-boundary protocols may well catch on.

—*Craig Partridge*
craig@tereschau.net

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**

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/

Fragments

JPNIC RPKI Guidelines Released

The *Japan Network Information Center* (JPNIC) has released a set of guidelines[1] aimed at mitigating unauthorized routing incidents on the Internet using *Resource Public Key Infrastructure* (RPKI) *Route Origin Authorizations* (ROA). These guidelines provide technical and operational recommendations to enhance the security and reliability of Internet routing. The objective of these guidelines is to promote the adoption of RPKI-based security measures. Targeting a broad audience that includes both managerial and engineering professionals in the ISP and network operations sectors, the document offers a structured approach to implementing and maintaining RPKI.

Developed with inputs from the *Japanese Network Operators Group* (JANOG), research from the Ministry of Internal Affairs and Communications cybersecurity initiatives, and expert consultations, the guidelines offer practical insights based on real-world deployment experiences.

The guidelines cover both organizational and technical aspects of RPKI implementation. They explain the business risks associated with unauthorized routes and highlight the importance of adopting RPKI to mitigate these threats. By understanding these risks, decision-makers can justify investment in RPKI and align their security strategies with industry best practices. For network operators, the guidelines offer step-by-step instructions on creating ROAs and deploying *Route Origin Validation* (ROV). These measures ensure that only legitimate route announcements are propagated, reducing the risk of route hijacking and improving overall network security.

The guidelines also outline role-based measures for different types of network operators. IP holders are required to create ROAs and maintain consistency between their ROA records and routing information to prevent discrepancies. *Autonomous System* (AS) operators are encouraged to implement ROV to filter out invalid routes, strengthening the security of the global routing system. The guidelines include real-world configuration examples for routers and outline security measures for BGP beyond RPKI, ensuring that operators have practical resources to facilitate implementation.

Version 1 of the guidelines is available now in Japanese (translatable) in web and PDF formats and is supplemented with practical configuration examples for ROV deployment on routers. JPNIC plans to update the guidelines regularly in collaboration with experts to incorporate evolving best practices and emerging threats and has also developed an online tool, *rov-check*[2], which allows network operators to verify whether their networks are effectively protected by ROV.

[1] **https://www.nic.ad.jp/ja/rpki/guideline/**

[2] **https://rov-check.nic.ad.jp/en**

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
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
Joseph Connolly
Steve Corbató
Brian Courtney
Beth and Steve Crocker
Dave Crocker
Kevin Croes
John Curran
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
Bill Fenner
Paul Ferguson
Ricardo Ferreira
Kent Fichtner
Ulrich N Fierz
Armin Fisslthaler
Michael Fiumano
The Flirble Organisation
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
Gulf Coast Shots
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
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
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
Martin Lillepuu
Roger Lindholm
Link Light Networks
Art de Llanos
Mike Lochocki
Chris and Janet Lonvick
Mario Lopez
Sergio Loreti
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
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
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
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
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
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
Phil Tweedie
Steve Ulrich
Unitek Engineering AG
John Urbanek
Martin Urwaleck
Bart Vanautgaerden
Betsy Vanderpool
Surendran Vangadasalam
Ramnath Vasudha
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
Jd Wegner
Tim Weil
Westmoreland Engineering Inc.
Rick Wesson
Peter Whimp
Russ White
Jurrien Wijlhuizen
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

JUNIPER NETWORKS

lacnic

LinkedIn

linx

NTT DATA

RIPE NCC RIPE NETWORK COORDINATION CENTRE

TEAM CYMRU

VERISIGN

WIDE PROJECT

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