

The Internet Protocol Journal

October 2024

Volume 27, Number 3

A Quarterly Technical Publication for
Internet and Intranet Professionals

FROM THE EDITOR

In This Issue

From the Editor	1
The Interop Shownet	2
“IPv6 Mostly” Experiment...	16
Fragments	22
Thank You.....	28
Call for Papers.....	30
Supporters and Sponsors.....	31

These days, it is common for Internet technology conferences to deploy a *temporary* network in convention centers or hotels to support Internet access for attendees and exhibitors. In some cases, these networks are used for special technology demonstrations. In this issue we will look at two examples of such networks.

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, or “Shownet,” 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. We will publish a separate article about the Japanese version of Shownet in a future edition. This time David Strom describes the history and evolution of the Interop Shownet. His article is dedicated to the memory of Daniel C. Lynch, the founder of Interop, who passed away earlier this year.

Work on the transition from IPv4 to IPv6 continues to be a major focus of several working groups in the *Internet Engineering Task Force* (IETF). As solutions are developed, technology events such as the *Asia Pacific Internet Conference on Operational Technologies* (APRICOT) provide end users an opportunity to experience IPv6 by simply selecting a designated Wi-Fi network on their devices. One such “IPv6 Mostly” experiment was conducted during APRICOT 2024 in Bangkok; Brian Candler describes it in our second article.

Ten years ago, *The Internet Protocol Journal* was relaunched with the help of numerous supporters, sponsors, and individual donors. Today we very much depend on this funding model, and once again we encourage you to make a donation or ask your organization to become a sponsor. We appreciate your feedback and suggestions. Please contact us via e-mail at: ipj@protocoljournal.org

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

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

ISSN 1944-1134

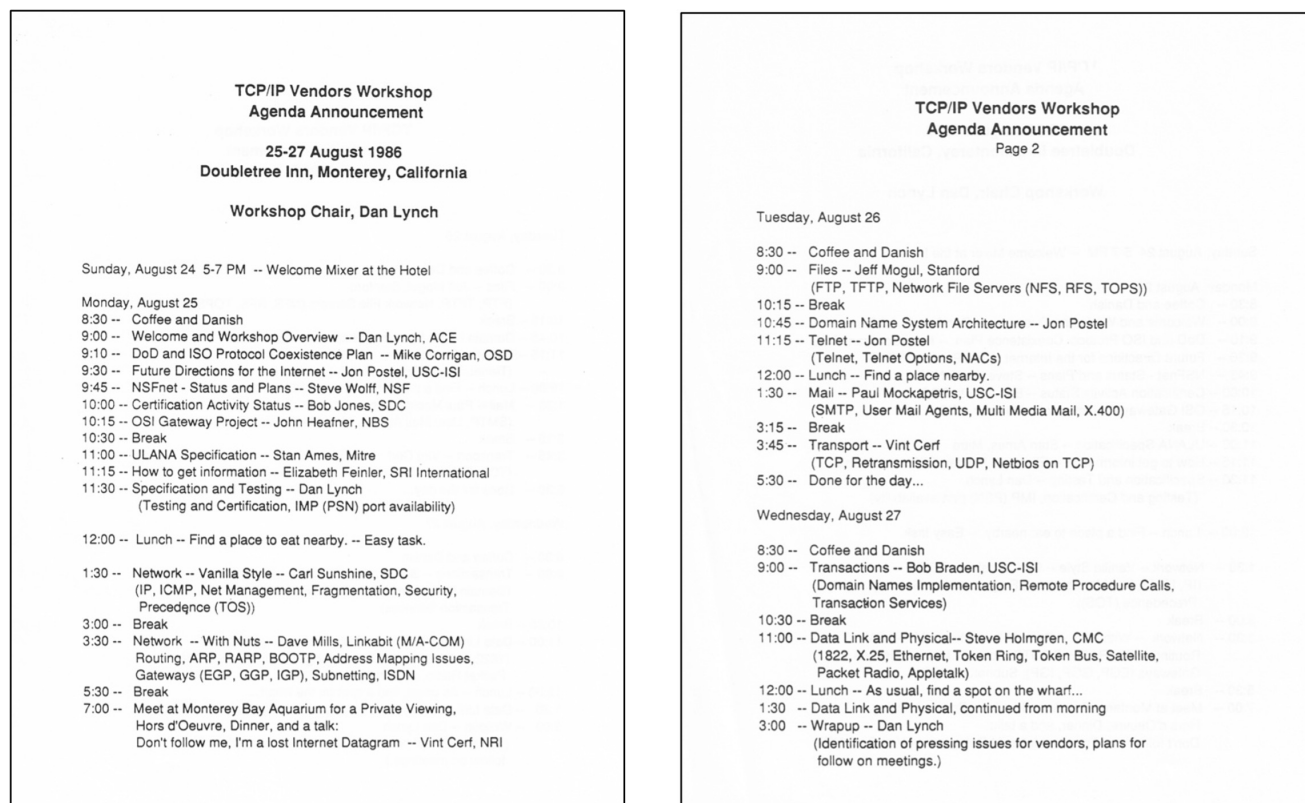
The Interop Shownet

by David Strom

This is the story about how a group of very dedicated people came together at the dawn of the Internet era to build something special, something unique and memorable. It was called the *Interop Show 'n Tel-Net*, later known as *InteropNet*, or *Shownet*, and it was created in September 1988 at the third *TCP/IP Interoperability Conference* held in Santa Clara, California. This story tells how it evolved and specifically how the larger context of this network became a powerful tool that moved the Internet from a mostly government-sponsored research project to a network that would support commercial businesses and could be used by millions of ordinary people in their daily lives. But before we consider what happened then, we must turn back the clock a couple of years.

In August 1986, a few very motivated people decided to teach others how to implement the early Internet protocols. This first conference, called the *TCP/IP Vendors Workshop*, was held in Monterey, California, and was by invitation only (See Figure 1).

Figure 1: TCP/IP Vendors Workshop Agenda, August 1986.


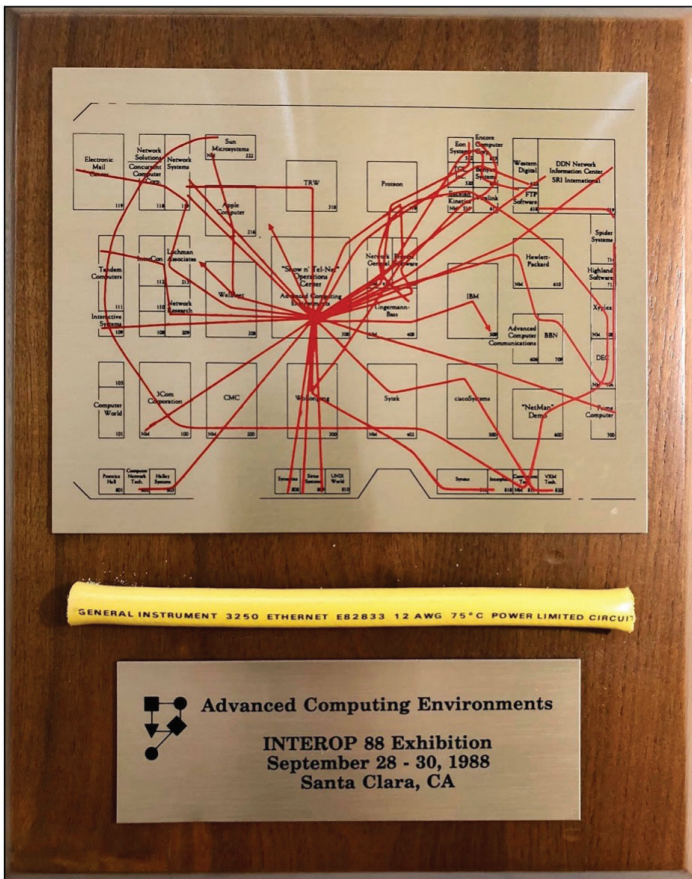


Speakers included Vint Cerf, who was at MCI at the time, Jon Postel, *Request For Comments* (RFC) Editor and at ARPANET before playing a key role in Internet administration, and Paul Mockapetris and Bob Braden, both at the *University of Southern California's Information Sciences Institute* (USC-ISI).

Two subsequent conferences were held the following March in Monterey and then December in Crystal City, Virginia, and both were called the *TCP/IP Interoperability Conference*. All three were unusual events for several reasons: first, the presenters and instructors were the actual engineers that developed the earliest Internet protocols. They were also there to impart knowledge, rather than sell products—mainly because few commercial products were yet invented. One of the instructors, Douglas Comer of Purdue University, wrote the first and best-selling book on the topic: *Internetworking with TCP/IP, Volume 1, Principles, Protocols and Architecture*.

By September of 1988, the format of the conference changed, and expanded beyond lectures to a more practical proving ground. The event was renamed once again, and so *Interop*—and its show network—was born. The mission was still to teach Internet technologies and protocols, but for the first time the event was used to test and demonstrate various Internet communications devices on an active computer network. That show used a variety of Ethernet cables to connect 51 exhibitors together, with T1 links to the NASA Ames Research Center in Mountain View, California, and the NSFNet in Ann Arbor, Michigan (Figure 2).

Figure 2: Left: Interop 88 Exhibition show network. Right: NETWORL+INTEROP 1996 advertisement.



THIS EVENT COULD REALLY GET ROUTY.

And the routing part is just for starters. After all, this is the world's top interoperability event.

At NetWorld+Interop® '96 Las Vegas, you'll get a chance to work with more than 50 hard-core networking experts like Wej, Doug, Robin and Steve. Together with other members of our Network Operations Center (NOC) team, they're in charge of operating our 6000 node, multi-vendor, multi-protocol network proving ground—the InteropNet™.

The NOC team designs and integrates classical networks with cutting-edge technologies.

So you get to see ATM running with existing networks. Or test 100base T vs. 100VG, routing vs. switching, client/server interoperability, Internet applications and more.

The InteropNet also connects more than 550 vendors with the latest networking solutions, like LAN emulation, Fibre Channel, 155Mbps Wireless, FDDI, 100VG-AwLan, ATM, frame relay transports...the list goes on and on.

Call today to get a FREE VIP Exhibition Pass! Or access the Web at <http://www.interop.com> to register and get complete event information.

NETWORL+INTEROP

Exhibition is April 2-4 • Conference is April 1-5 • Las Vegas Convention Center • 800-488-2883

Interop is a registered trademark of INTEROP Corp. InteropNet is a service mark of Interop Inc. All other names are the property of their respective holders.

The Interop conference quickly grew into a worldwide series of events^[1] with multiple shows held in different cities that were attended by tens of thousands of visitors with more than a thousand connected booths. In those early years, the largest shows were held in Tokyo, which began in 1994 and continued annually (with a pause because of the pandemic), with the latest show held in mid-2024. This year's show spanned over 500 vendors' booths and drew about 40,000 visitors each of its three days. The Tokyo Interop is also where the *ShowNet* (this is the chosen capitalization for the Tokyo show) not only has survived, but also has thrived, and continues to innovate and demonstrate Internet interoperability to this day. Many products had their world or Japanese debuts at various Tokyo Interop events, including Cisco's XDR and 8608 Router and NTT's Open APN.

My Own Interop Journey with Network Computing Magazine

Before I discuss the evolution of Interop and the role and history of its show network, I should first mention my own personal journey with Interop. In 1990, I was in the process of creating the first issue of *Network Computing* magazine for CMP Media. Our first issue was going to debut at the Interop show, the second time it was held that year in the San Jose, California, Convention Center.

The publisher and I both thought this place was the best one for our debut for several reasons. First, our magazine was designed for similar motivations—to demonstrate what worked in the new field of computer networking. We had designed our publication around a series of laboratories that had the same equipment found in a typical corporate office, including wide-area links and a mixture of PC MS-DOS, Apple Macintosh, and Unix devices and even a DEC minicomputer connected together. Second, we wanted to make a “big splash,” and our salespeople were already showing prototype issues ahead of the show to entice advertisers to sign up. Finally, *Network Computing's* booth would be connected to the Shownet and the greater Internet, just like many of the exhibitors who were trying out some product for the very first time.

One other feature about *Network Computing* that set us apart from other business trade magazines at the time: each bylined article would contain the email address of the author, so that readers could contact them with questions and comments. I wanted to use the domain **cmp.com** and set up an actual Internet presence, but alas I was overruled by management, so we ended up using a gateway maintained by one of the departments at UCLA where a couple of our editors were housed. While posting an author's email contact is now common, it was a radical notion at the time.

That 1990 Interop conference began my own personal journey of numerous shows around the world through the years, including speaking and teaching, as well as covering them for various business and networking publications that I would write for.

The Earliest Days of Interop

The Interop show in the late 1980s was a markedly different trade show from others of its era. At the time, trade shows with networked booths were non-existent. By way of perspective, up until that point in those early years, there were two kinds of conferences: one focused on the trade show with high-priced show floors and fancy exhibits. There, exhibitors were forced to “pay to play,” meaning if they bought booth space, they could secure a speaking slot at the associated conference. The other was a more staid affair that was a gathering of the engineers and actual implementers. Interop was a notable early example bridging the two: it looked like a trade show but was more of a conference, all in the guise of getting better commercial products out into the marketplace. It helped that it had its roots in those early TCP/IP conferences.

“You could see the Internet in a room, thanks to the Shownet, with hundreds of nodes talking to each other. That was unique for its time,” said Carl Malamud. “The Shownet was the most complex Internet installation you could do at any moment of time.”

Malamud would play several key roles in the development of various early Internet-based projects, including running the first Internet-based radio station, and he was a Shownet volunteer in 1991. In addition, Interop commissioned him to write the book *Exploring the Internet: A Technical Travelogue*.^[6]

That complexity has been true from the moment the Shownet was first conceived to the present day. Many of the Internet protocols—both in their earliest years and up to the present era—were debugged over the Shownet: volunteers recall testing NetBIOS over TCP/IP, 10BaseT Ethernet, SNA over TCP/IP, *Fiber Distributed Data Interface* (FDDI), *Simple Network Management Protocol* (SNMP), *Internet Protocol Version 6* (IPv6), various versions of segment routing, and numerous others. That extensive protocol catalog is a testament to the influence and effectiveness of the Shownet, and how enduring a concept it has been over the course of Internet history. Steve Hultquist, who was part of the early *Network Operations Center* (NOC) teams, remembers that the first version of 3Com’s 100BaseT switch—with “serial number 1”—was installed on the show network.

The force behind Interop was Dan Lynch, who passed away earlier this year. Lynch foresaw the commercial Internet and designed Interop to hasten its adoption. He based Interop on a series of efforts to bring together TCP/IP vendors, and the proto-Interop shows that were run in the middle 1980s that were more “Plugfests” or “Connectathons,” where vendors would try out their products. The main difference was those efforts deployed mostly proprietary protocols, whereas Interop ran on open source.

He told Sharon Fisher in November 1987: “There are millions of PCs out there and they’re starting to get networked in meaningful ways, not just in little printer-sharing networks.”

Part of his vision—and those that he recruited—was the notion of interoperability that could be used as a selling point and as an alternative to single-vendor proprietary networks from IBM, Digital Equipment Corporation, and others that were common in that era. Larry Lang, who worked for many years at Cisco, said, “The reassurance that it was okay to give up having ‘one throat to choke’ came from confidence that the equipment was interoperable. It is hard to remember a time when that was a worry, but it surely was.”

Part of Lynch’s vision was to ensure that proving interoperability was a very simple litmus test: did the product being exhibited work as advertised in a real-world situation? The answer to that question seems like common sense, but doing so in a trade show context was a relatively rare idea. And while it was a simple question, the answer was usually anything but simple, and sometimes the reasons why some product didn’t work—or didn’t work all the time—was what made the Shownet a powerful product improvement tool. That is just as true today as it was back then. The more realistic the Shownet, the more often it would expose these special circumstances that would bring out the bugs and other implementation problems.

It would prove to be a potent and enduring vision.

What Does Interoperability Mean?

The notion of interoperability seems so common sense now—and indeed it is the default position for most of the current networking world. However, in the early days of the Internet it was fraught with problems in terms of both larger-scale implementations and smaller issues that would prevent products from working reliably. One of Sharon Fisher’s articles in *Computerworld* in 1991^[2] speaks about TCP/IP this way: “The astounding thing is not how gracefully it performs but that it performs at all. TCP/IP is not for everyone.” Times certainly have changed in the 33 years since that was written. Today the notion of Internet connectivity, using TCP/IP protocols, is a given assumption in any computing product—from smart watches to the largest mainframe computers.

Those early implementation differences plagued both large and small vendors alike, and required a meeting of the minds where the protocol specifications weren’t exacting enough to ensure its success, or where bugs took time to resolve. Enter the Interop conference. As a reminder, in those early days the popular IP applications were based on the *File Transfer Protocol* (FTP), the *Simple Mail Transfer Protocol* (SMTP), and the *Simple Network Management Protocol* (SNMP). The web was still being invented and far away from being the de facto smash hit that it is today. Video conferencing and streaming didn’t exist. Telephones still ran on non-IP networks.

One of the early casualties was the *Open Systems Interconnection* (OSI) series of protocols promulgated by the *International Organization for Standardization* (ISO).

It was precisely because of the interoperability among TCP/IP products and “the failure of OSI to effectively demonstrate interoperability in the early 1990s that was the final nail in its coffin,” said Brian Lloyd, who worked at Telebit at the time. There are other stories of the defeat of OSI, such as this one in the *IEEE Spectrum*.^[3]

The Relationship Among the Shownet, the Conference Tutorials, and the NOC

To accomplish Lynch’s vision, Interop was not only the Shownet, but also its interaction with two other elements that became force multipliers in the quest for interoperability. These elements were the tutorials that were given before the opening of the show floor, and the NOC team that ran the network itself. All three had an important synergy to promote the actual practice of interoperability among different vendors’ products: not only in the demonstration of what worked with what but also in the discovery of protocol mismatches or programming errors so that new equipment could be made to interoperate.

Dave Crocker, who authored many RFCs and served on the Interop program committee that selected speakers in the 1990s, called out this tripartite structure of Shownet, NOC, and conference as a major strength of Interop. “Interop was able to contrast the technologies of the Internet with the interoperability of non-Internet technologies, such as IBM’s *Systems Network Architecture* (SNA). It had very pragmatic implications and wasn’t just promoting marketing speak.”

Many of the engineers who developed those early protocols and techniques and other pieces of Internet technology taught the tutorials (and as I mentioned, I taught a few of them during those early years, in addition to serving with Crocker on the program committee), so that others could learn how to best implement them. Here is where Interop contained its secret sauce: the people who taught the tutorials were the people who contributed to the underlying protocols and code, in some cases code so new it was changing over the duration of the show itself. “It was only after getting to Interop that we found out how few options were actually used by most implementations, and only then did we have access to the larger Internet and various versions of Unix computers,” said Brian Lloyd. “It was real bleeding edge stuff back then and the place to go for product testing and see how marketers and engineers would work together.”

And the NOC was a real one, like what could be found at large corporations, monitoring the network for anomalies and using it to debug various implementations leading up to the opening moments of the show. “It was unusual for its time,” said Fisher in another article in *Infoworld*.^[4] “The NOC team was infamous in the trade press for its tours and the time members took to explain things to us,” she said. Malamud recalls that the NOC had a strict “no suits” policy, meaning that its denizens were engineers that rolled up their sleeves and got stuff done.

All of this happened with very few paid staff: most of the people behind the Shownet and NOC were volunteers who came back, show after show, to work on setting things up and then taking them down after the show ended. That was, and to some extent still is, a very high-pressure environment: imagine wiring up a large convention center and connecting all of its conference rooms with a variety of network cabling. Several of the original Interop Shownet and NOC volunteers are continuing the tradition by helping to build and run the *Internet Engineering Task Force* (IETF) event networks at every IETF meeting around the world.

One of the more infamous moments of Interop was the *Internet Toaster*, created originally by John Romkey for the 1990 show^[5].

“I wanted to get people thinking of SNMP not just as getting variables, but for control applications, a wider vision. So we had an SNMP controlled toaster. If you put bread in the toaster, and set a variable in SNMP, the toaster would start toasting. A whole *Management Information Base* (MIB) was written for it, including how you wanted the toast, and whether it was a bagel or Wonder Bread. I ended up with lots and lots of bread in my garage. It got a lot of attention, but I don’t think that managing your kitchen through SNMP is very practical today.”

Dave Buerger, who was an early tech journalist at CMP, remembers Interop as having “a strange sense of awe unfolding for everyone as we glimpsed the possibilities of global connectivity. Exhibits on the show floor were more experiments in connecting their booths to the rest of the world.”

Construction of the Earliest Shownets (before 1993)

To say that Lynch was very persuasive is perhaps a big understatement. He convinced people who were quirky, unruly, or difficult to work with to spend lots of time pulling things together. “Dan allowed us to do stuff that the usual convention wouldn’t normally allow, and managed people that weren’t used to being managed,” said Malamud.

Peter de Vries was one of the earliest Shownet builders when he was working at Wollongong as one of the early Internet vendors. He ended up working for Interop for three years before opening the West Coast office for FTP Software. He remembers Lynch “dragging people kicking and screaming into using the Internet” back in the late 1980s and early 1990s. “But he was a fun guy to work for, and he had an unusual management style where he didn’t issue demands but convinced you to do something through more subtle suggestions, so by the time you did it you were convinced you had the original idea.”

These volunteers would essentially be working year-round, especially once the calendar was filled with multiple shows per year. Back in those early years, the convention centers didn’t care about cabling, and hadn’t yet figured out that having a more permanent physical networking plant could be used as an asset for attracting future meetings.

“We were often the first show to hang cables from the ceiling, and it wasn’t easy to do,” said Malamud, who chronicled the 1991 Shownet assembly^[6]. The first Interop shows used thick Ethernet cables that required a great deal of finesse to work with; de Vries recalls they had to pass wires through expansion joints and other existing holes in the walls and floors, wires that didn’t easily bend around corners.

“Each network tap took at least ten minutes of careful drilling to attach to this thick cable.” He has many fond memories of Lynch: “My goal was to try to get everyone to use TCP/IP, but Dan took it to the next step and showed that TCP could be a useful tool, something better than a fax. He was a real visionary.”

Ethernet—in all of its variations over the years, including early implementations of 10BaseT and 100-megabit speeds—wasn’t the only cabling choice for Interop; the show would expand to fiber and *Token Ring* cabling as part of its mission. Brian Chee was one of those volunteers who remembers having to re-terminate 150 different fiber strands across the high catwalks of the convention spaces. “We even had to terminate the fiber on the roof of the convention center to connect it to the Las Vegas Hilton across the street,” he said.

Getting all that cabling up in the air wasn’t an easy task either. Patrick Mahan worked on the San Francisco show in 1992 and recalls that he and other networking volunteers were paired with the union electrical workers on a series of scissor lifts. “We needed multiple lifts operating in tandem to raise them, and we had just started hanging a 100-foot length [cable] when a loud air-horn goes off and each lift immediately starts descending to the floor, because it was time for the mandatory union 15-minute break. It took about three minutes before the cable bundles started breaking apart and crashing to the cement floor. You could hear the glass in the fiber cables breaking!”

The Vegas climate made installations difficult, especially when its non-air-conditioned convention halls reached temperatures of 110 degrees Fahrenheit outside. “Many convention centers don’t turn on their air conditioning until the night before the show begins, so it was a particularly harsh environment,” said Glenn Evans, who worked as both a volunteer and an employee of Interop during the late 1990s and early 2000s. “Vegas in May is very dry, and static electricity is a big issue. We fried several switch ports inadvertently and spent long nights adding static filters to avoid it because some shows had more than 1,200 connections across their networks.” Evans emphasized that the install teams relied on “redneck engineering to come up with creative solutions, and it didn’t have to be perfect, [it] just had to work for five days.”

The cabling had to be laid out three times for each Shownet. The volunteers would have access to the convention center for a day months before any actual show. They would lay out the first cable segments and add connectors, then roll them up and store them in a warehouse. Then before the show there would be a “hot staging” event where the cables were connected to their equipment racks and tested.

Finally, several nights before the show began saw the real deployment at the convention center, which would span several 24-hour days before the actual opening. Those long nights were epic: de Vries recalls falling asleep in the middle of one night at the top of a 15-foot ladder, only to be gently awakened by the only other person in the convention center at the time. “Those installations nearly killed me!”

Many of the participants during those early years were motivated by a sense of common purpose, that their efforts were directly contributing to the Internet and its usefulness. “I loved that we could help the overall industry get stuff right,” said Hultquist. “They were some of the smartest people that I have ever worked with and were constantly pushing the envelope to try to deploy all sorts of emerging technologies.”

But the physical plant was just one issue; once the cabling was in place, the real world of getting equipment up and running across these networks was challenging. In those early Interops, equipment was often at the cutting edge, and engineers would make daily or even minute-by-minute changes to their protocol stacks and application code.

James van Bokkelen was the president of FTP Software then, and he recalls seeing the Shownet in 1988 crash while running BSD v4.3, thanks to a buggy version of one TCP/IP command. Turns out the bug was present in Cisco’s routers that were used on the Shownet. “It took a few minutes of scampering before everything was in place and we got Shownet back online,” he said. Scampering indeed: the volunteers had to compare notes, debug their code, and reboot equipment often located at different ends of the convention floor.

“We were getting alpha software releases during the show. This network created an environment where people had to fix things in real time in real production environments,” said Hultquist. “Wellfleet, 3Com, and Cisco were all sending us router firmware updates so their gear could interoperate with each other. I loved that we could help the overall industry get stuff right.”

At one of the 1991 Interop events, “FDDI completely melted down,” said Merike Kaeo, who at the time was working for Cisco in charge of their booth and volunteering in the NOC. “There was some obscure bug where a router reboot wasn’t enough, you had to reset the FDDI interface adapter separately. It didn’t take all that long to get things running, thankfully.”

Some of the problems were far more mundane, such as using equipment with NiCad batteries that had very short shelf life. Chee recalls that one Fluke engineering director got tired of trying to get these batteries replaced with Lithium-ion batteries. “He would send his team up to the rafters with network test equipment that had very short battery life; they were quickly replaced in their newer products.”

As more Shownets were brought up over the years, they had built-in redundant—and segregated—links. “We all played a part to make sure that after 1991, we would have a stable portion that would run reliably and put any untested equipment on another network that wouldn’t bring the working network down after the show started,” said Bill Kelly, who worked for Cisco in its early days.

de Vries said the first couple of Interop Shownets had less than 10 miles of cabling, which grew by 1991, according to Malamud, to having more than 35 miles of cabling, connecting a series of ribs, each one running down an aisle of the convention floor or some other well-defined geographic area. “Each rib had both Ethernet and Token Ring connected to an equipment rack with various routers,” said Malamud.

“There were two backbones that connected 50 different subnets, one based on FDDI and the other on Ethernet, which in turn were connected via T-1 lines to NASA Ames Research Center and Bay Area Internet points.”

Bill Kelly, who worked on the Shownet NOC while he was at Cisco, developed a three-stage model that covered a product lifecycle. “The first stage is using the IETF RFCs to try to make something work. Then the second stage is when a vendor is late to market and must figure out how to play nicely with the incumbents and the standards. The third stage is mostly commodity products, and everything works as advertised.”

The Middle Years (1994–1999)

The Internet—and Interop—were both growing quickly during this time. New Internet protocols and RFCs were being created frequently, and applications—and dot-com businesses—sprang up without any business plans, let alone initial paying customers. There were new venues each year in Europe, shows in Sydney, Australia, and Singapore, and Sao Paulo, Brazil. Some years had as many as seven or eight different shows, each with its own Shownet that needed to be customized for the exhibit halls in these cities.

Let’s return to 1986 for a moment. That year Novell began its own trade shows, called *NetWorld*, to explain its growing *Netware* community. By 1994, these shows had grown, and that is when Novell and Interop merged their shows, calling them *NetWorld+Interop*. This moniker held until 2004, when a series of technical media companies purchased Interop.

“Shownet didn’t change much after the Novell merger. We could accommodate their stuff at the edge, but it didn’t impact the core network,” said Hultquist. For the Shownet team, Netware was just another protocol to interoperate across. Despite Novell’s influence, during these years, TCP/IP became a networking standard. So did the cabling that made up the Shownet: “In the mid-1990s, a lot of the cable plant could be reused from show to show, with a standard set of 29-strand multimodal fiber with quick connectors and 48 strands of Category 5 copper cable for the ribs,” said Evans.

TCP/IP evolved too: by the end of the 1990s, the protocols and Ethernet hardware became commodities and were both factory-installed in millions of endpoint devices. “The Internet was becoming more standardized, and Interop became less of an experiment and more of a technology demonstration,” said Evans.

Nevertheless, vendors tried to differentiate themselves with quirky exhibits, pushing the envelope of connectivity. One stunt happened during the 1995 Interop at the Broadcom booth, which demonstrated Ethernet signals over barbed wire. “The wires were ugly and rusty and had nasty little barbs all over them,” according to one description written years later.^[6,11]

By 1999, the Shownet split into two separate parts: the live production network connecting the exhibit booths called *InteropNet*, and *InteropNet Labs* used for showcasing new technologies and products. Back then, these new technologies included VoIP, VPNs, and other “hot technologies,” according to a post by Tim Greene on CNN.^[7] Several market forces caused this situation. First, more and more conferences began promoting the idea of Internet connectivity for both attendees and vendor participants. “As that reality dawned on people, the Interop Shownet became an increasingly useless anachronism,” said Larry Lang, who was part of the team building Cisco’s support for FDDI at that time. “As our competition became Wellfleet rather than IBM, why would we want to participate in an expensive and time-consuming display that suggested complete equivalence among all the products?”

Hultquist was quoted in that CNN piece saying that attendees “won’t know whether a piece of equipment really worked because of the demands placed on them by more experimental or untested products.”

A second issue had to do with striking a balance between established vendors and newcomers. Kelly remembers the relationship between Cisco and Interop to be “complicated because we were the market leader and if we just donated equipment without any technical support, we ran the risk of outsiders misconfiguring the devices. Interop was also used to dealing with small engineering groups and not pesky marketing types that wanted to know the value of participating in the show.” Plus, long-running contributors to the original Shownets often got a jump on developing new gear and interacting with products that weren’t yet on the market.

By the end of the 1990s, the Shownet staging operation had also split into two. Prior to that moment, each Shownet would be staged in a Silicon Valley warehouse. But then the show runners for Tokyo decided to set up their own facilities to stage and constitute their own Shownet, and reformulated their NOC team from local talent, where they continue to build and demonstrate interoperability to the present day.

The New Millennium of Interop

Interop continued to grow in the new millennium. Two notable events affected the Shownet. The day the towers fell in New York in 2001 was also the day that the fall Interop show in Atlanta started. Many of the Shownet volunteers recall how quickly their network became the main delivery of news and video feeds to those attendees who were stuck in Atlanta, since all flights were grounded for the next several days. Brian Chee remembers that “within minutes of the disaster we maxed out the twin OC-12 WAN connections into the Shownet. We brought up streaming video of *CNN Headline News* over IP multicast, and that cut our wide-area traffic substantially, while at the same time it was an impressive demonstration of that technology.”

But then a few years later another event happened. “The day the Slammer virus hit, in 2003, we had just gone into production across the Shownet. That virus hurt our network throughput just enough that all our monitoring devices were useless,” said Chee. “But the NOC team was able to characterize the problem within a few minutes, and we were able to use air gapped consoles to reset routers and filter out the virus-infected packets.” That is as real world as it gets and is an example of how the Shownet proved its worth, time and time again.

But what is amazing is how enduring the legacy of Shownet continues to be. For example, during the 2019 Tokyo Interop, it played a critical role in demonstrating the interoperability among various segment routing vendors running over IPv6, resulting in a draft Internet document.^[8] I had an opportunity to review a draft report from the Tokyo team about the 2024 network that will be published in a future IPJ issue. “We faced varied challenges and considerations to achieve this while serving user traffic,” they wrote^[9].

Subsequent Tokyo shows—indeed the now sole survivors of the Interop legacy—would continue to draw on a huge talent pool of local talent. This year’s show had more than 650 volunteer engineers, including 30 alone to operate its NOC. “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,” said Takashi Tomine, who was part of the NOC team. The NOC occupies an impressive amount of show-floor real estate, where it continues to serve as a teaching and demonstration tool, as well as a working network nerve center.

It also is an opportunity for university students and junior staff to obtain hands-on experience in network operations and spend two weeks touching technology in ways that they might not have in their jobs or classrooms. The NOC team conducts walking tours, wherein guides describe what these teams have done in the many Interops held elsewhere down through the years. The 2024 show endures in another way, the “hot staging” model that was developed more than 30 years ago at the first Silicon Valley shows. The team has a total of eight days to assemble the network, and a few hours to take it apart after the show ends.

“First, we install every device in the right place on the racks, turn on the devices, and check their status. Checking device statuses 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,” said Tomine.

His article will be published in a subsequent IPJ issue that goes into further details about the Tokyo show and how it grew over the years.

Figure 3: Interop Tokyo
2024 ShowNet Walking Tour



The 2024 Tokyo Interop showcased several new technologies, or technologies used in new and innovative ways. For example, the Shownet shared streaming video content with three geographically distributed TV broadcasting stations, all over IP networks. The team built a special media operations center to control these broadcasts and to demonstrate real-time video recording and editing of several conference sessions and demonstrations. In that respect, it was back to the future when the first multicast IP streams were broadcast years ago.

“The coolest thing we got out of working at Interop is that technology doesn’t happen without the people, and the people involved were some of the hardest-working and smartest people that you’ll ever meet. They checked their egos at the door, and solved problems jointly,” said Evans. “It was run like a democratic dictatorship, where everyone had a say.”

Acknowledgements

This article wouldn’t be possible without the help of numerous volunteers and staff at Interop events of the past, including Bill Alderson, Karl Auberbach, Dave Buerger, Brian Chee, Dave Crocker, Peter J.L. de Vries, Glenn Evans, Sharon Fisher, Connie Fleenor, Steve Hultquist, Merike Kaeo, Bill Kelly, Larry Lang, Brian Lloyd, Carl Malamud, Jim Martin, Naoki Matsuhira, Ryota Roy Motobayashi, Takashi Tomine, and James van Bokkelen.



Dedication

This article is dedicated to the memory of Daniel Courtney Lynch, August 16, 1941 – March 30, 2024, founder of the Interop events and whose vision gave us the Interop Shownet.

References and Further Reading

- [1] Ryota Motobayashi, “Report on INTEROP,”
<http://motobayashi.net/interop/>
- [2] Sharon Fisher, “TCP/IP: Effective, yes. Well supported, definitely. But if you’re looking for a graceful way to link PCs to hosts, you’d better wait a while,” *ComputerWorld*, October 7, 1991.
- [3] Andrew L. Russell, OSI: The Internet That Wasn’t,” *IEEE Spectrum*, July 29, 2013.
<https://spectrum.ieee.org/osi-the-internet-that-wasnt>
- [4] Sharon Fisher, “Products Help Users Manage TCP/IP-Based Networks,” *Infoworld*, October 10, 1988.
- [5] John Romkey, as told to Bernard Aboba, “How PC-IP Came to Be,” *The Internet Archive*, August 7, 2011.
- [6] Carl Malamud, “San Jose,” a chapter from *Exploring the Internet: A Technical Travelogue*, ISBN-13, 978-0132968980, January 1992.
- [7] Tim Greene, “Bleeding edge tech booted from show net,” CNN, May 11, 1999.
- [8] 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.
- [9] Takashi Tomine, Ryo Nakamura, and Ryota Motobayashi, “ShowNet at Interop Tokyo: A Continuously Evolving Demonstration Network,” to be published in a future edition of *The Internet Protocol Journal*.
- [10] Karl Auerbach, “Our Scrapbook (The Ghost of Interop Past)”:
<https://cavebear.com/archive/interop/>
- [11] Dr. Howard Johnson, “So Good, It Works on Barbed Wire,” *EDN Magazine*, July 5, 2001.
- [12] “Building the INTEROPnet 93,” YouTube video.

DAVID STROM has been writing about technology since 1986 for various print and online publications, including *PC WeekleWeek*, *ComputerWorld*, *Infoworld*, *eeTimes*, *Tom’s Hardware*, *CSOonline*, *SiliconANGLE.com*, *Dark-Reading.com*, and dozens of others. He last wrote in IPJ Volume 23, No. 2 about his experiences selling his own class C block of IPv4 addresses back in 2020. He was the founding editor-in-chief of *Network Computing*, and he helped build numerous technical and editorial websites focusing on networking, security, and communications. He has written two books, *Internet Messaging* (1998, Prentice Hall) with Marshall T. Rose and *Home Networking Survival Guide* (2001, Osborne/McGraw Hill). He has two daughters and lives in St. Louis and can be found online at strom.com

Experimental IPv6-only Network at APRICOT 2024

by Brian Candler, NSRC

The *Asia Pacific Regional Internet Conference on Operational Technologies* (APRICOT)^[0] is the Internet Network Operators Summit for the Asia Pacific region. Every year, a wireless network conference is deployed to provide connectivity for hundreds of delegates, with a separate “IPv6-only” SSID as an alternative for people to try. This year, we decided to experiment with a new approach, by using some of the recent mechanisms designed for “IPv6-mostly” networks to build a better “IPv6-only” network.

What Is IPv6-mostly?

“IPv6-mostly” is a way to gracefully sunset *Internet Protocol Version 4* (IPv4) on dual-stack access networks. The work in this area has been driven in part by Google’s enterprise network, which has become so large that they ran out of RFC 1918^[1] private addresses. Technically, there are two pieces to this plan:

- A new “IPv6-only preferred” option for *Dynamic Host Configuration Protocol Version 4* (DHCPv4) (option 108, RFC 8925)^[2]. By requesting this option, a client declares that it is willing to run in a single-stack, IPv6-only mode. And by returning this option, the DHCPv4 server confirms that the network is happy to work this way too. The client then doesn’t configure itself with any IPv4 address.
- A new “PREF64” *Neighbor Discovery Option for Router Advertisements* (RFC 8781)^[3]. This option tells the client that a *Network Address Translator 64* (NAT64) is available, and what prefix to use.

If both of these conditions are true, the client configures itself with a *Customer-Side NAT46 Translator* (CLAT), with a hidden private IPv4 address. Any IPv4 application traffic is routed through this translator and carried across the network as IPv6 until it reaches the *Provider-Side NAT64* (PLAT), where it is converted back to IPv4. This whole mechanism is called *464XLAT*.

The end result is that you can interact with IPv4 resources—even using IPv4 literals, like `ping 8.8.8.8`—when running on an IPv6-only network. In effect, your IPv6 network doubles as a large block of private addresses behind a NAT. A big advantage of this approach is that there is no need to use DNS64 to generate fake AAAA records for IPv4-only destinations.

“IPv6-mostly” is supported by modern versions of macOS (13+), iOS, and Android. Any other clients will simply continue with regular dual-stack operation, but overall the usage of your DHCPv4 address pools will decrease.

Using IPv6-mostly Features for IPv6-only

For APRICOT, we wanted to build a pure IPv6-only network, not dual-stack. But we also wanted to enable the CLAT in client devices that support it in order to get maximum compatibility with IPv4. Here is how we enabled it: the pieces were all built inside an Ubuntu 22.04 *Virtual Machine* (VM) running on a compact *Next Unit of Computing* (NUC) computer. The NUC is a line of small-form-factor barebone computer kits designed by Intel.

First, we needed a DHCPv4 server that would respond to clients that requested option 108, granting them permission to run IPv6-only. Regular DHCP servers like *ISC DHCP* and *Kea DHCP* are quite happy to do that. However, we also did *not* want to respond to clients who *didn't* support option 108; if we did, we'd have to offer them an IPv4 address, and we'd be back to a dual-stack network.

I couldn't find an off-the-shelf DHCPv4 server that was capable of working this way, so I found a modular DHCP server in Go called *coredhcp*^[4] and created a new plugin^[5] to implement the desired behavior. This plugin has now been merged into the main codebase.

Second, I needed to send router advertisements with the PREF64 option. The conference routers were Arista Layer 3 switches, and although they have this feature in very recent firmware, it wasn't available in the version we were using.

Therefore, I used Linux's *radvd*^[6] to perform the router advertisements. This feature was not available in the latest released version, only git HEAD, so I had to compile *radvd* from source. Since it's not possible for one router to send advertisements on behalf of another, it meant that the VM where *radvd* was running also had to act as the gateway for the IPv6-only network, turning the VM into a router for IPv6 traffic.

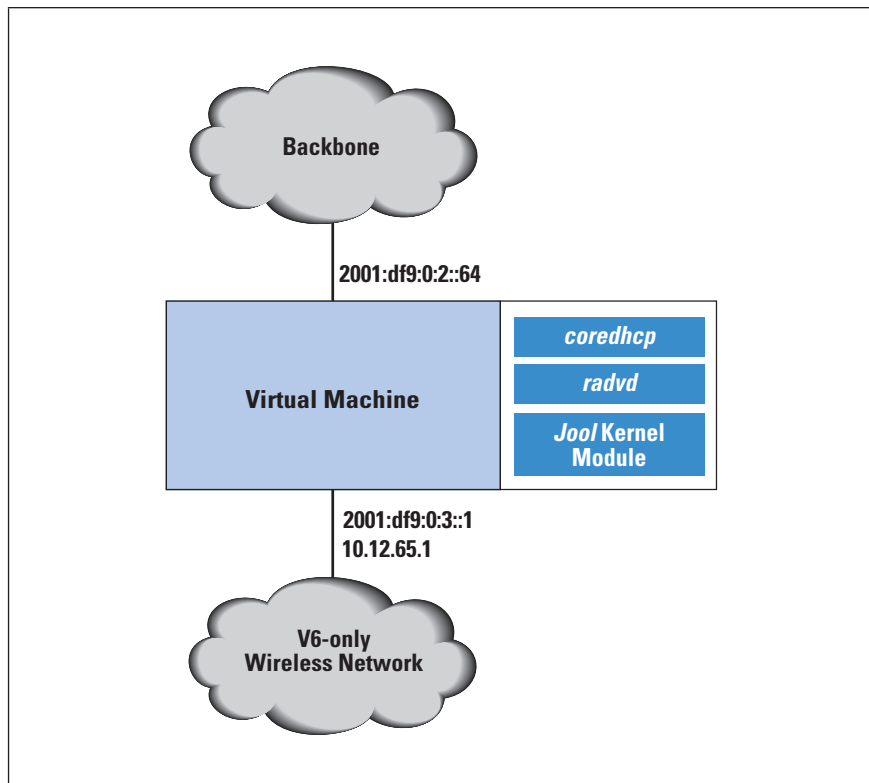
And finally, I needed a NAT64 translator somewhere on the network. This process could potentially have been done on a conference router, but since the IPv6 traffic had to pass through the VM anyway, I decided to implement NAT64 on the VM as well. I implemented it using *Jool*,^[7] a kernel module for NAT64 (refer to Figure 1 on the following page).

Configuration

We bundled all the components into a single VM running Ubuntu Linux 22.04. The VM acted as the router for the IPv6-only network, forwarding IPv6 packets in the kernel. It had an IPv6 address on the conference backbone, and an IPv4 address on a separate /64 subnet for the wireless network, with a static route on the upstream conference router.

We also configured the VM with an IPv4 address on the wireless-facing interface, but only so that the DHCPv4 server had an address to bind to. The VM offered no IPv4 addresses to clients, and enabled no IPv4 routing or NAT.

Figure 1: APRICOT 2024 Network Configuration for “IPv6 Mostly”



The other components running on the VM follow:

- The *Jool* kernel module, to perform the NAT64 translation;
- The *radvd* daemon, to announce the wireless prefix and the NAT64 prefix to clients; and
- The *coredhcp* server, to return DHCPv4 option 108 to those clients that requested it. It was also configured to act as a stateless DHCPv6 server, to give out DNS server addresses and domain search lists for any clients that didn't support Router Advertisement options *Recursive DNS Server Option* (RDNSS) and *Domain Name System Search List* (DNSSL).

Problems

The basic reachability to the IPv4 Internet via NAT64 seemed to work well. However, we were plagued by clients being repeatedly removed from the IPv6-only network by the Cisco *Wireless LAN Controller* (WLC), at seemingly random intervals. Strangely, this didn't seem to affect clients on the main conference SSID on the same access points. It turns out that several underlying issues were to blame.

First, the client private CLAT address **192.0.0.2** was leaking out as the source IP address on various multicast packets (such as multicast DNS, and Chrome doing service discovery on UDP port 1900). Logs showed that the wireless controller had a feature called “IP Theft or Reuse” that would add client MAC addresses to an exclusion list if it saw the same source IP address from multiple clients. We were able to turn that option off.

Second, some clients were being kicked off regularly every 10 minutes. The Cisco WLC had a setting where the wireless access points would perform dynamic channel reassignment every 10 minutes. We increased it to 24 hours.

Finally, when the conference was nearly over, we discovered another per-SSID setting, “IPv4 DHCP required,” which we also turned off. We believe that fixed the remaining problems.

There was one other major issue: clients would lose the ability to reach IPv4 destinations for a few minutes at a time, without being kicked off the wireless network—IPv6 connectivity continued to work. Using *tcpdump*, we saw that IPv6 neighbor discovery on the VM was forgetting about the IPv6 address of the CLAT. It turns out that the version of Jool in the Ubuntu 22.04 package repositories is old (v4.1.7), and this problem is known^[8]. It was straightforward to upgrade the package to the latest release, v4.1.11.

Apart from these problems, the CLAT/NAT64 mechanism worked very well for those devices that supported it. Remaining issues were minor: *traceroute* to an IPv4 destination showed only “*” for every hop, and the macOS ssh client didn’t work when given the “-4” flag (although it did work with an IPv4 literal address). Otherwise, it was just like being on a dual-stack network.

Conference Usage and Compatibility

From DHCP logs, I found that 142 unique devices had attempted to use the IPv6-only SSID, and of those, 115 (81%) supported DHCP option 108. That’s a surprisingly high proportion, representing a high usage of Apple laptops, iOS phones, and Android phones amongst delegates. Those devices should have gotten a good experience from the network, if it weren’t for the wireless disconnection issues.

The other 27 devices would have had a much worse experience. They got no DHCPv4 response, so they retried repeatedly, configured themselves with an IPv4 link-local address (**169.254.x.x**), and still would have been able to reach only Internet sites with IPv6 addresses.

We could have improved compatibility for these clients somewhat by providing a DNS64 service, which fakes AAAA records for DNS names that have only A records. However, these DNS settings would have applied to all hosts on the network, meaning that even those clients supporting option 108 would also have been exposed to fake DNS responses. I thought that the experiment was more useful without it, because the NAT64/DNS64 combination is already well-known and tested.

Post-conference Updates

Since the conference, I learned that it may no longer be necessary to provide any DHCPv4 server. Recent versions of MacOS will enable the CLAT even if the device has no IPv4 address at all.

However, an absent DHCPv4 server causes clients to keep sending DHCPDISCOVER messages, generating unnecessary load on both the clients and the network. Returning DHCPv4 option 108 stops the clients from doing this constant resending. There is also an older standard, “Auto-Config” DHCP option 116, in RFC 2563^[9], but DHCP logs from the IPv6-only conference network showed no clients using this option, so it appears to be obsolete.

I also discovered a problem about the choice of the NAT64 prefix. A *Well-Known Prefix* (WKP), **64:ff9b::/64**, is defined in RFC 6052^[10]. But if you use it, you will find that clients will be unable to connect via NAT64 to private addresses such as RFC 1918, because that is mandated by RFC 6052 section 3.1. If you want to use NAT64 on a typical home or enterprise network, and still be able to reach internal devices on private addresses, you will need to avoid the WKP. The conference network used a *Unique Local Address* (ULA) prefix (**fd64::/64**) instead.

Conclusion

IPv6-only using the IPv6-mostly mechanisms works surprisingly well, and is only going to improve over time as Windows^[11] and Linux add support for it.

Personally, I’d be quite happy to run this way at home, except that my Mikrotik router has no NAT64 capability. (RouterOS versions 7.8 and later do have the PREF64 router advertisement option though)^[12].

References and Further Reading

- [0] Asia Pacific Regional Internet Conference on Operational Technologies: <https://apricot.net/>
- [1] Yakov Rekhter, Robert G. Moskowitz, Daniel Karrenberg, Geert Jan de Groot, and Eliot Lear, “Address Allocation for Private Internets,” RFC 1918, February 1996.
- [2] Lorenzo Colitti, Jen Linkova, Michael C. Richardson, and Tomek Mrugalski, “IPv6-Only Preferred Option for DHCPv4,” RFC 8925, October 2020.
- [3] Lorenzo Colitti and Jen Linkova, “Discovering PREF64 in Router Advertisements,” RFC 8781, April 2020.
- [4] *coredhcp* server: <https://github.com/coredhcp/coredhcp>
- [5] Plugin for *coredhcp*:
<https://github.com/coredhcp/coredhcp>. Option to Disable Stateless Auto-Configuration in IPv4 Clients,” RFC 2563, May 1999.
- [6] *radvd*: <https://github.com/radvd-project/radvd>
- [7] *Jool* kernel module: <https://github.com/NICMx/Jool>
- [8] “Jool when translating always drops the packet due to an error,” <https://github.com/NICMx/Jool/issues/382>
- [9] Ryan Troll, “DHCP Option to Disable Stateless Auto-Configuration in IPv4 Clients,” RFC 2563, May 1999.

- [10] Congxiao Bao, Christian Huitema, Marcelo Bagnulo, Mohamed Boucadair, and Xing Li, “IPv6 Addressing of IPv4/IPv6 Translators,” RFC 6052, October 2010.
- [11] Tommy Jensen, “Windows 11 Plans to Expand CLAT Support,” *Microsoft Tech Community Networking Blog*, March 7, 2024.
- [12] Mikrotik change log.
- [13] The video of my presentation at APRICOT 2024.
- [14] The slides from my presentation at APRICOT 2024.
- [15] Ondřej Caletka, “Deploying IPv6-mostly access networks,” RIPE Labs.
- [16] Jen Linkova, “Mission Possible: Turning off IPv4 in Google Enterprise Network,” presented at RIPE 87,

BRIAN CANDLER is a graduate of the University of Cambridge, UK. Having originally started in microelectronics design, he has developed systems and networking at several large UK-based ISPs. He delivers training courses for the Network Startup Resource Center in developing regions of the world, covering topics such as campus network design and network monitoring and management, and develops a virtual training platform used to provide hands-on lab exercises. He can be reached at: brian@nsrc.org

Brian Candler at APRICOT 2024



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

IAB Workshop on Barriers to Internet Access of Services (BIAS)

The *Internet Architecture Board* (IAB) organizes workshops about topics of interest to the community that bring diverse experts together, raise awareness, and possibly identify the next steps that can be explored by the community. The IAB held its *Barriers for Internet Access of Services* (BIAS) fully online workshop during the week of January 15, 2024.

The Internet is a crucial component of our critical infrastructure that wields a significant influence on various aspects of society. It serves as a vital tool for advancing the United Nations' *Sustainable Development Goals* (SDGs)^[1] and upholding human rights on a global scale. Thus the absence of meaningful access to digital infrastructure and services amounts to a form of disenfranchisement. The barriers to meaningful access to Internet-based services and applications are increasing, posing challenges that persist even when Internet connectivity is available, thereby resulting in unequal information and service access.

The workshop solicited position papers about barriers to accessing content and services on the Internet, for example, based on filtering, blocking as well as due to general inequality of technological capabilities, like device or protocol limitations. 19 position papers were submitted to the workshop of which 12 papers were selected for publication^[2]. Two invited talks were also presented based on published papers. There were 40 participants in the workshop over three days.

This marked my first IAB workshop since joining the board. I am delighted to have collaborated with Mirja Kühlewind, Mallory Knodel, Tommy Pauly, and Christopher A. Wood in organizing this event. The themes of censorship, circumvention techniques, and the digital divide have surfaced in various IAB discussions lately. Our goal for this workshop was to present reports, expert opinions, and ignite discussions on these topics. Through this experience, I gained valuable insights and strongly believe that the IETF community must remain mindful of these crucial issues when designing protocols. It is imperative to ensure that we create the most secure, user-friendly protocols for all Internet users.

This article provides a short overview of the workshop discussion. However, if you would like to learn more you can also check out the initial draft version of the IAB workshop report^[3], or watch the entire thing on YouTube^[4]. The workshop was organized into three main themes across three days based on the submitted papers.

Community Networks

Community Networks are self-organized networks which are wholly owned by the community and thus provide an alternative mechanism to bring connectivity and Internet services to those places that lack commercial interest. Discussion ranged from highlighting the need for measuring *Quality of Experience* (QoE) for Community Networks, to the potential role a *Content Delivery Network* (CDN) can play in Community Networks, to the role of satellite networks, and finally, to the vital role of the spectrum in this space.

Digital Divide The digital divide refers to disparities in access to the Internet and services. It signifies the gap between those who have effective and meaningful access to digital technologies and those who do not. Discussion recognized three key aspects of the digital divide: differences between population demographics in the provision of online resources by governments, inequality in the use of multilingual domains and email addresses, and increased costs for end-user downloads of contemporary websites' sizes. There was a general recognition that there may be more technical aspects of the digital divide that were not presented.

Censorship Censorship is the legal control or suppression of what can be accessed, published, or viewed on the Internet. This discussion focused on reports of censorship as observed during recent years in different parts of the world, as well as on the use of and expectation for censorship circumvention tools, mainly the use of secure VPN services. This included censorship reports from India and Russia, where censorship has changed significantly recently, highlighting the legal frameworks and court acts that put obligations on regional network providers to block traffic. Further, measurements to validate the blocking, as well as analyses of how blocking is implemented were also discussed.

Next Steps The discussion highlighted the need for the technical community to address the management gaps and document best practices for Community Networks including listing of manageability considerations explicitly for Community Networks. Further, the need to build consensus on solutions that have the most significant impact in fostering digital inclusion and the need to further promote them was discussed. We need to continue to work towards enhancing our protocols ensuring user privacy, develop further protocols that enable more transparency on filtering and new VPN-like services. Further discussion of these topics could happen in the *Global Access to the Internet for All (GAIA)*^[5], *Human Rights Protocol Considerations (HRPC)*^[6], *Privacy Enhancements and Assessments (PEARG)*^[7], and *Measurement and Analysis for Protocols (MAPRG)*^[8] research groups, based on the relevance to each group.

—Druv Dhody, IAB Member
dd@dhruvdhody.com

References

- [1] United Nations Department of Economic and Social Affairs, “Sustainable Development, The 17 Goals.”
- [2] IAB Workshop on Barriers to Internet Access of Services (BIAS) (biasws) Materials.
- [3] IAB Barriers to Internet Access of Services (BIAS) Workshop Report.
- [4] YouTube recordings from BIAS Workshop.
- [5] Global Access to the Internet for All (GAIA) Research Group.
- [6] Human Rights Protocol Considerations Research Group.
- [7] Privacy Enhancements and Assessments Research Group
- [8] Measurement and Analysis for Protocols Research Group.

IAB Statement on the Risks of Attestation on the Open Internet

While attestation of client software and hardware is a useful tool for preventing abuse or fraud on the Internet, the use of such attestation as a barrier to access otherwise open protocols and services would negatively impact the evolution of the Internet as a whole.

Openness and the empowerment of end users are core values of the IETF. RFC 3935^[1], Section 4.1, explains this as part of the IETF's mission statement:

“We want the Internet to be useful for communities that share our commitment to openness and fairness. We embrace technical concepts such as decentralized control, edge-user empowerment and sharing of resources, because those concepts resonate with the core values of the IETF community. These concepts have little to do with the technology that's possible, and much to do with the technology that we choose to create.”

The Internet is built upon the idea that anyone who implements the appropriate standards should be able to interoperate on the Internet. Many of the core services that run on the Internet, such as email and the web, are designed to be openly accessible in this way. Adding client attestation into otherwise open systems can significantly reduce openness for the Internet broadly. A recent “Web Environment Integrity”^[2] proposal has highlighted this risk, although such models pose a risk beyond just the web.

Attestation of client software and hardware is distinct from user authentication. User authentication verifies the identity of a user or a credential associated with a user, and is compatible with any implementation that supports the correct form of authentication. In contrast, attestation of client software and hardware places explicit restrictions on the implementations that are allowed to participate in the protocol. For services that have intentionally restricted access, such client attestation (as described in *Remote ATtestation procedureS* (RATS), RFC 9334^[3]) is a valuable security measure, particularly when used in conjunction with user authentication. However, this approach is not appropriate for openly accessible services.

Allowing clients to use a variety of software as long as it is protocol-compliant is an essential part of the IETF development process and the openness of the Internet. Although customized or open-source software can also be used to circumvent client-side security measures, the continuing viability of open software is required for continued innovation. Restricting access via attestation of software or hardware would limit the development of new protocols and extensions to existing protocols, lock users into a limited ecosystem of applications, and hamper the ability to audit implementations, conduct measurements, or perform essential security research.

If client attestation signals are used in open services to mitigate fraud or abuse, they should be designed to only signal the authenticity of a user or client without imposing strict software or hardware requirements. They should also be designed such that attestation is not required, but has a clear backup behavior when attestation is not possible. IETF-based protocols such as *Privacy Pass* [RFC 9576] attempt to provide a protocol that can be deployed in ways that promote user privacy without exposing detailed identifiers about the client systems that are being used. Fundamentally, attesting specific properties about a networking client (for example, there is some human user involved in this interaction) maintains the openness of the Internet, whereas attesting that a specific piece of software is in use does not and should be avoided.

The IAB invites those in the industry and standards community working on client attestation in open services to engage with the relevant IETF working groups (in particular, Privacy Pass^[4] and RATS^[5]), and encourages those groups to focus on defining safe deployment models for attestation and abuse prevention that will not put the openness of the Internet at risk.

References

- [1] Harald Alvestrand, “A Mission Statement for the IETF,” RFC 3935, October 2004.
- [2] Web Environment Integrity API:
<https://github.com/explainers-by-googlers/Web-Environment-Integrity/>
- [3] Henk Birkholz, Dave Thaler, Michael Richardson, Ned Smith, and Wei Pan, “Remote ATtestation procedureS (RATS) Architecture,” RFC 9334, January 2023.
- [4] Privacy Pass Working Group:
<https://datatracker.ietf.org/wg/privacypass/about/>
- [5] Remote ATtestation ProcedureS Working Group:
<https://datatracker.ietf.org/wg/rats/about/>

Internet Community Encouraged to Submit Event Proposals for UA Day 2025

Event proposals are now being accepted for the third annual *Universal Acceptance (UA) Day*, to be held between 1 March and 30 May 2025. UA Day, held annually, is an opportunity to rally local, national, regional, and global communities and organizations to spread UA awareness and to promote UA adoption with key stakeholders.

UA is a technical best practice that ensures all valid domain names and email addresses, regardless of script, language or character length, can be equally used by all Internet-enabled applications, devices, and systems. Co-organized by the *Universal Acceptance Steering Group* (UASG) and the *Internet Corporation for Assigned Names and Numbers* (ICANN), UA Day 2025 will consist of various virtual and in-person events held by the UASG, ICANN, global partners, and regional and local organizations.

UA Day 2025 builds on the success of *UA Day 2024* and the inaugural *UA Day* in 2023. Together, the 2023 and 2024 events attracted approximately 15,000 participants worldwide across dozens of countries. These milestone events have helped mobilize technical and language communities, companies, governments and *Domain Name System* (DNS) industry stakeholders to champion UA on a global scale.

Those interested in organizing a UA Day event must complete the *UA Day Event Proposal Form*^[1] by 11 October 2024. ICANN will provide limited support for proposed UA Day events based on this group's recommendations. The following types of events are eligible for support:

- *UA Awareness*: Provide a high-level overview of UA and *Email Address Internationalization* (EAI), the benefits of being UA-ready, basic technical concepts related to UA and next steps for becoming UA-ready.
- *UA Technical Training*: Provide in-depth training on becoming EAI-ready for email system administrators and on becoming UA-ready for software developers.
- *UA Academic Curricula*: Work with academic faculty members and experts to integrate *Internationalized Domain Names* (IDNs) and UA-related topics into existing technical curricula and design a roadmap.
- *UA Adoption*: Conduct a UA adoption exercise and share challenges and solutions to becoming UA-ready, and document your experience. Please note that advance preparatory work and review are required in order to qualify for a UA Adoption event.
- *UA Regional Strategy*: Discuss appropriate mechanisms for promoting UA adoption at the local, regional, and national levels.

Proposals will be considered from all relevant organizations, including international, regional and local organizations, technology organizations and companies, open-source communities, standards bodies, email service providers, academia, industry groups, and language communities.

UA is considered a foundational requirement for the continued expansion of the Internet. Since 2009, the landscape for domain names has changed markedly—in overall number of *Top-Level Domain Names* (TLDs) available, TLD character length and scripts available. However, the checks used by many software applications to validate domain names and email addresses often use rules that do not fully support Universal Acceptance. Achieving UA ensures everybody has the ability to experience the full social and economic power of the Internet using their chosen domain name and email address that best aligns with their interests, business, culture, language, and script.

Questions can be directed to UAProgram@icann.org. A full UA Day event calendar will be published in due course. In the meantime, interact with the UASG on social media (X, Facebook and LinkedIn) using the hashtag **#Internet4All**.

The UASG is a community-led initiative that was formed in 2015 and funded by ICANN. It consists of volunteers from many companies, governments, and community groups. The UASG works to raise awareness of the importance of UA globally, provide free resources to organizations to help them become UA-ready, and measure the progress of UA adoption. To learn more, visit <https://uasg.tech/>

ICANN's mission is to help ensure a stable, secure, and unified global Internet. To reach another person on the Internet, you need to type an address—a name or a number—into your computer or other device. That address must be unique so computers know where to find each other. ICANN helps coordinate and support these unique identifiers across the world. ICANN was formed in 1998 as a nonprofit public benefit corporation with a community of participants from all over the world. To learn more, visit <https://icann.org>

References

- [1] UA Day Event Proposal Form:
<https://tinyurl.com/UADayForm>

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.

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

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

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

The Internet Protocol Journal is published under the “CC BY-NC-ND” Creative Commons Licence. Quotation with attribution encouraged.

This publication is distributed on an “as-is” basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.

Supporters and Sponsors

<p><i>Supporters</i></p>  	<p><i>Diamond Sponsors</i></p> <p>Your logo here!</p>
<p><i>Ruby Sponsors</i></p>  	<p><i>Sapphire Sponsors</i></p> <p>Your logo here!</p>

Emerald Sponsors

























Corporate Subscriptions












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

Ole J. Jacobsen, Editor and Publisher

Editorial Advisory Board

Dr. Vint Cerf, VP and Chief Internet Evangelist
Google Inc, USA

John Crain, Senior Vice President and Chief Technology Officer
Internet Corporation for Assigned Names and Numbers

Dr. Steve Crocker, CEO and Co-Founder
Shinkuro, Inc.

Dr. Jon Crowcroft, Marconi Professor of Communications Systems
University of Cambridge, England

Geoff Huston, Chief Scientist
Asia Pacific Network Information Centre, Australia

Dr. Cullen Jennings, Cisco Fellow
Cisco Systems, Inc.

Merike Kaeo, Founder and vCISO
Double Shot Security

Olaf Kolkman, Principal – Internet Technology, Policy, and Advocacy
The Internet Society

Dr. Jun Murai, Founder, WIDE Project
Distinguished Professor, Keio University
Co-Director, Keio University Cyber Civilization Research Center, Japan

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

The title "The Internet Protocol Journal" is a trademark of Cisco Systems, Inc. and/or its affiliates ("Cisco"), used under license. All other trademarks mentioned in this document or website are the property of their respective owners.

Printed in the USA on recycled paper.

