

How the Internet Came to Be

By Vincent Cerf

The birth of the ARPANET

My involvement began when I was at UCLA doing graduate work from 1967 to 1972. There were several people at UCLA at the time studying under Jerry Estrin, and among them was Stephen Crocker. Stephen was an old high-school friend, and when he found out that I wanted to do graduate work in computer science, he invited me to interview at UCLA.

When I started graduate school, I was originally looking at multiprocessor hardware and software. Then a Request For Proposal came in from the Defense Advanced Research Projects Agency, DARPA. The proposal was about packet switching, and it went along with the packet-switching network that DARPA was building.

Several UCLA faculty were interested in the RFP. Leonard Kleinrock had come to UCLA from MIT, and he brought with him his interest in that kind of communications environment. His thesis was titled Communication Networks: Stochastic Flow and Delay, and he was one of the earliest queuing theorists to examine what packet-switch networking might be like. As a result, the UCLA people proposed to DARPA to organize and run a Network Measurement Center for the ARPANET project.

This is how I wound up working at the Network Measurement Center on the implementation of a set of tools for observing the behavior of the fledgling ARPANET. The team included Stephen Crocker; Jon Postel, who has been the RFC editor from the beginning; Robert Braden, who was working at the UCLA computer center; Michael Wingfield, who built the first interface to the Internet for the Xerox Data System Sigma 7 computer, which had originally been the Scientific Data Systems (SDS) Sigma 7; and David Crocker, who became one of the central figures in electronic mail standards for the ARPANET and the Internet. Mike Wingfield built the BBN 1822 interface for the Sigma 7, running at 400 Kbps, which was pretty fast at the time.

Around Labor Day in 1969, BBN delivered an Interface Message Processor (IMP) to UCLA that was based on a Honeywell DDP 516, and when they turned it on, it just started running. It was hooked by 50 Kbps circuits to two other sites (SRI and UCSB) in the four-node network: UCLA, Stanford Research Institute (SRI), UC Santa Barbara (UCSB), and the University of Utah in Salt Lake City.

We used that network as our first target for studies of network congestion. It was shortly after that I met the person who had done a great deal of the architecture: Robert Kahn, who was at BBN, having gone there from MIT. Bob came out to UCLA to kick the tires of the system in the long haul environment, and we struck up a very productive collaboration. He would ask for software to do something, I would program it overnight, and we would do the tests.

One of the many interesting things about the ARPANET packet switches is that they were heavily instrumented in software, and additional programs could be installed remotely from BBN for targeted data sampling. Just as you use trigger signals with oscilloscopes, the IMPs could trigger collection of data if you got into a certain state. You could mark packets and when they went through an IMP that was programmed appropriately, the data would go to the Network Measurement Center.

There were many times when we would crash the network trying to stress it, where it exhibited behavior that Bob Kahn had expected, but that others didn't think could

happen. One such behavior was reassembly lock-up. Unless you were careful about how you allocated memory, you could have a bunch of partially assembled messages but no room left to reassemble them, in which case it locked up. People didn't believe it could happen statistically, but it did. There were a bunch of cases like that.

My interest in networking was strongly influenced by my time at the Network Measurement Center at UCLA.

Meanwhile, Larry Roberts had gone from Lincoln Labs to DARPA, where he was in charge of the Information Processing Techniques Office. He was concerned that after building this network, we could do something with it. So out of UCLA came an initiative to design protocols for hosts, which Steve Crocker led.

In April 1969, Steve issued the very first Request For Comment. He observed that we were just graduate students at the time and so had no authority. So we had to find a way to document what we were doing without acting like we were imposing anything on anyone. He came up with the RFC methodology to say, "Please comment on this, and tell us what you think."

Initially, progress was sluggish in getting the protocols designed and built and deployed. By 1971 there were about nineteen nodes in the initially planned ARPANET, with thirty different university sites that ARPA was funding. Things went slowly because there was an incredible array of machines that needed interface hardware and network software. We had Tenex systems at BBN running on DEC-10s, but there were also PDP8s, PDP-11s, IBM 360s, Multics, Honeywell... you name it. So you had to implement the protocols on each of these different architectures. In late 1971, Larry Roberts at DARPA decided that people needed serious motivation to get things going. In October 1972 there was to be an International Conference on Computer Communications, so Larry asked Bob Kahn at BBN to organize a public demonstration of the ARPANET.

It took Bob about a year to get everybody far enough along to demonstrate a bunch of applications on the ARPANET. The idea was that we would install a packet switch and a Terminal Interface Processor or TIP in the basement of the Washington Hilton Hotel, and actually let the public come in and use the ARPANET, running applications all over the U.S.

A set of people who are legendary in networking history were involved in getting that demonstration set up. Bob Metcalfe was responsible for the documentation; Ken Pogran who, with David Clark and Noel Chiappa, was instrumental in developing an early ring-based local area network and gateway, which became Proteon products, narrated the slide show; Crocker and Postel were there. Jack Haverty, who later became chief network architect of Oracle and was an MIT undergraduate, was there with a holster full of tools. Frank Heart who led the BBN project; David Walden; Alex McKenzie; Severo Ornstein; and others from BBN who had developed the IMP and TIP.

The demo was a roaring success, much to the surprise of the people at AT&T who were skeptical about whether it would work. At that conference a collection of people convened: Donald Davies from the UK, National Physical Laboratory, who had been doing work on packet switching concurrent with DARPA; Remi Despres who was involved with the French Reseau Communication par Paquet (RCP) and later Transpac, their commercial X.25 network; Larry Roberts and Barry Wessler, both of whom later joined and led BBN's Telenet; Gesualdo LeMoli, an Italian network researcher; Kjell Samuelson from the Swedish Royal Institute; John Wedlake from British Telecom; Peter Kirstein from University College London; Louis Pouzin who led the

Cyclades/Cigale packet network research program at the Institute Recherche d'Informatique et d'Automatique (IRIA, now INRIA, in France). Roger Scantlebury from NPL with Donald Davies may also have been in attendance. Alex McKenzie from BBN almost certainly was there.

I'm sure I have left out some and possibly misremembered others. There were a lot of other people, at least thirty, all of whom had come to this conference because of a serious academic or business interest in networking.

At the conference we formed the International Network Working Group or INWG. Stephen Crocker, who by now was at DARPA after leaving UCLA, didn't think he had time to organize the INWG, so he proposed that I do it.

I organized and chaired INWG for the first four years, at which time it was affiliated with the International Federation of Information Processing (IFIP). Alex Curran, who was president of BNR, Inc., a research laboratory of Bell Northern Research in Palo Alto, California, was the U.S. representative to IFIP Technical Committee 6. He shepherded the transformation of the INWG into the first working group of 6, working group 6.1 (IFIP WG 6.1).

In November 1972, I took up an assistant professorship post in computer science and electrical engineering at Stanford. I was one of the first Stanford acquisitions who had an interest in computer networking. Shortly after I got to Stanford, Bob Kahn told me about a project he had going with SRI International, BBN, and Collins Radio, a packet radio project. This was to get a mobile networking environment going. There was also work on a packet satellite system, which was a consequence of work that had been done at the University of Hawaii, based on the ALOHA-Net, done by Norman Abramson, Frank Kuo, and Richard Binder. It was one of the first uses of multiaccess channels. Bob Metcalfe used that idea in designing Ethernet before founding 3COM to commercialize it.

The birth of the Internet

Bob Kahn described the packet radio and satellite systems, and the internet problem, which was to get host computers to communicate across multiple packet networks without knowing the network technology underneath. As a way of informally exploring this problem, I ran a series of seminars at Stanford attended by students and visitors. The students included Carl Sunshine, who is now at Aerospace Corporation running a laboratory and specializing in the area of protocol proof of correctness; Richard Karp, who wrote the first TCP code and is now president of ISDN technologies in Palo Alto. There was Judy Estrin, a founder of Bridge Communications, which merged with 3COM, and is now an officer at Network Computing Devices (NCD), which makes X display terminals. Yogen Dalal, who edited the December 1974 first TCP specification, did his thesis work with this group, and went on to work at PARC where he was one of the key designers of the Xerox Protocols. Jim Mathis, who was involved in the software of the small-scale LSI-11 implementations of the Internet protocols, went on to SRI International and then to Apple where he did MacTCP. Darryl Rubin went on to become one of the vice presidents of Microsoft. Ron Crane handled hardware in my Stanford lab and went on to key positions at Apple. John Shoch went on to become assistant to the president of Xerox and later ran their System Development Division. Bob Metcalfe attended some of the seminars as well. Gerard Lelann was visiting from IRIA and the Cyclades/Cigale project, and has gone on to do work in distributed computing. We had Dag Belsnes from University of Oslo who did work on the correctness of protocol design; Kuninobu Tanno (from Tohoku University); and Jim Warren, who went on to found the West Coast Computer Faire. Thinking about

computer networking problems has had a powerful influence on careers; many of these people have gone on to make major contributions.

The very earliest work on the TCP protocols was done at three places. The initial design work was done in my lab at Stanford. The first draft came out in the fall of 1973 for review by INWG at a meeting at University of Sussex (September 1973). A paper by Bob Kahn and me appeared in May 1974 in IEEE Transactions on Communications and the first specification of the TCP protocol was published as an Internet Experiment Note in December 1974. We began doing concurrent implementations at Stanford, BBN, and University College London. So effort at developing the Internet protocols was international from the beginning. In July 1975, the ARPANET was transferred by DARPA to the Defense Communications Agency (now the Defense Information Systems Agency) as an operational network.

About this time, military security concerns became more critical and this brought Steve Kent from BBN and Ray McFarland from DoD more deeply into the picture, along with Steve Walker, then at DARPA.

At BBN there were two other people: William Plummer and Ray Tomlinson. It was Ray who discovered that our first design lacked and needed a three-way handshake in order to distinguish the start of a new TCP connection from old random duplicate packets that showed up later from an earlier exchange. At University College London, the person in charge was Peter Kirstein. Peter had a lot of graduate and undergraduate students working in the area, using a PDP-9 machine to do the early work. They were at the far end of a satellite link to England.

Even at the beginning of this work we were faced with using satellite communications technology as well as ARPANET and packet radio. We went through four iterations of the TCP suite, the last of which came out in 1978.

The earliest demonstration of the triple network Internet was in July 1977. We had several people involved. In order to link a mobile packet radio in the Bay Area, Jim Mathis was driving a van on the San Francisco Bayshore Freeway with a packet radio system running on an LSI-11. This was connected to a gateway developed by Virginia Strazisar at BBN. Ginny was monitoring the gateway and had artificially adjusted the routing in the system. It went over the Atlantic via a point-to-point satellite link to Norway and down to London, by land line, and then back through the Atlantic Packet Satellite network (SATNET) through a Single Channel Per Carrier (SCPC) system, which had ground stations in Etam, West Virginia, Goonhilly Downs England, and Tanum, Sweden. The German and Italian sites of SATNET hadn't been hooked in yet. Ginny was responsible for gateways from packet radio to ARPANET, and from ARPANET to SATNET. Traffic passed from the mobile unit on the Packet Radio network across the ARPANET over an internal point-to-point satellite link to University College London, and then back through the SATNET into the ARPANET again, and then across the ARPANET to the USC Information Sciences Institute to one of their DEC KA-10 (ISIC) machines. So what we were simulating was someone in a mobile battlefield environment going across a continental network, then across an intercontinental satellite network, and then back into a wireline network to a major computing resource in national headquarters. Since the Defense Department was paying for this, we were looking for demonstrations that would translate to militarily interesting scenarios. So the packets were traveling 94,000 miles round trip, as opposed to what would have been an 800-mile round trip directly on the ARPANET. We didn't lose a bit!

After that exciting demonstration, we worked very hard on finalizing the protocols. In the original design we didn't distinguish between TCP and IP; there was just TCP. In

the mid-1970s, experiments were being conducted to encode voice through a packet switch, but in order to do that we had to compress the voice severely from 64 Kbps to 1800 bps. If you really worked hard to deliver every packet, to keep the voice playing out without a break, you had to put lots and lots of buffering in the system to allow sequenced reassembly after retransmissions, and you got a very unresponsive system. So Danny Cohen at ISI, who was doing a lot of work on packet voice, argued that we should find a way to deliver packets without requiring reliability. He argued it wasn't useful to retransmit a voice packet end to end. It was worse to suffer a delay of retransmission.

That line of reasoning led to separation of TCP, which guaranteed reliable delivery, from IP. So the User Datagram Protocol (UDP) was created as the user-accessible way of using IP. And that's how the voice protocols work today, via UDP.

Late in 1978 or so, the operational military started to get interested in Internet technology. In 1979 we deployed packet radio systems at Fort Bragg, and they were used in field exercises. The satellite systems were further extended to include ground stations in Italy and Germany. Internet work continued in building more implementations of TCP/IP for systems that weren't covered. While still at DARPA, I formed an Internet Configuration Control Board chaired by David Clark from MIT to assist DARPA in the planning and execution of the evolution of the TCP/IP protocol suite. This group included many of the leading researchers who contributed to the TCP/IP development and was later transformed by my successor at DARPA, Barry Leiner, into the Internet Activities Board (and is now the Internet Architecture Board of the Internet Society). In 1980, it was decided that TCP/IP would be the preferred military protocols.

In 1982 it was decided that all the systems on the ARPANET would convert over from NCP to TCP/IP. A clever enforcement mechanism was used to encourage this. We used a Link Level Protocol on the ARPANET; NCP packets used one set of one channel numbers and TCP/IP packets used another set. So it was possible to have the ARPANET turn off NCP by rejecting packets sent on those specific channel numbers. This was used to convince people that we were serious in moving from NCP to TCP/IP. In the middle of 1982, we turned off the ability of the network to transmit NCP for one day. This caused a lot of hubbub unless you happened to be running TCP/IP. It wasn't completely convincing that we were serious, so toward the middle of fall we turned off NCP for two days; then on January 1, 1983, it was turned off permanently. The guy who handled a good deal of the logistics for this was Dan Lynch; he was computer center director of USC ISI at the time. He undertook the onerous task of scheduling, planning, and testing to get people up and running on TCP/IP. As many people know, Lynch went on to found INTEROP, which has become the premier trade show for presenting Internet technology.

In the same period there was also an intense effort to get implementations to work correctly. Jon Postel engaged in a series of Bake Offs, where implementers would shoot kamikaze packets at each other. Recently, FTP Software has reinstated Bake Offs to ensure interoperability among modern vendor products.

This takes us up to 1983. 1983 to 1985 was a consolidation period. Internet protocols were being more widely implemented. In 1981, 3COM had come out with UNET, which was a UNIX TCP/IP product running on Ethernet. The significant growth in Internet products didn't come until 1985 or so, where we started seeing UNIX and local area networks joining up. DARPA had invested time and energy to get BBN to build a UNIX implementation of TCP/IP and wanted that ported into the Berkeley UNIX release in v4.2. Once that happened, vendors such as Sun started using BSD as the base of commercial products.

The Internet takes off

By the mid-1980s there was a significant market for Internet-based products. In the 1990s we started to see commercial services showing up, a direct consequence of the NSFNet initiative, which started in 1986 as a 56 Kbps network based on LSI-11s with software developed by David Mills, who was at the University of Delaware. Mills called his NSFNet nodes "Fuzzballs."

The NSFNet, which was originally designed to hook supercomputers together, was quickly outstripped by demand and was overhauled for T1. IBM, Merit, and MCI did this, with IBM developing the router software. Len Bozack was the Stanford student who started Cisco Systems. His first client: Hewlett-Packard. Meanwhile Proteon had gotten started, and a number of other routing vendors had emerged. Despite having built the first gateways (now called routers), BBN didn't believe there was a market for routers, so they didn't go into competition with Wellfleet, ACC, Bridge, 3COM, Cisco, and others. The exponential growth of the Internet began in 1986 with the NSFNet. When the NCP to TCP transition occurred in 1983 there were only a couple of hundred computers on the network. As of January 1993 there are over 1.3 million computers in the system. There were only a handful of networks back in 1983; now there are over 10,000.

In 1988 I made a conscious decision to pursue connection of the Internet to commercial electronic mail carriers. It wasn't clear that this would be acceptable from the standpoint of federal policy, but I thought that it was important to begin exploring the question. By 1990, an experimental mail relay was running at the Corporation for National Research Initiatives (CNRI) linking MCI Mail with the Internet. In the intervening two years, most commercial email carriers in the U.S. are linked to Internet and many others around the world are following suit.

In this same time period, commercial Internet service providers emerged from the collection of intermediate-level networks inspired and sponsored by the National Science Foundation as part of its NSFNet initiatives. Performance Systems International (PSI) was one of the first, spinning off from NYSERNet. UUNET Technologies formed Altnet; Advanced Network and Systems (ANS) was formed by IBM, MERIT, and MCI (with its ANS CO+RE commercial subsidiary); CERFNet was initiated by General Atomics which also runs the San Diego Supercomputer Center; JVNNet became GES, Inc., offering commercial services; Sprint formed Sprintlink; Infonet offered Infolan service; the Swedish PTT offered SWIPNET, and comparable services were offered in the UK and Finland. The Commercial Internet eXchange was organized by commercial Internet service providers as a traffic transfer point for unrestricted service.

In 1990 a conscious effort was made to link in commercial and nonprofit information service providers, and this has also turned out to be useful. Among others, Dow Jones, Telebase, Dialog, CARL, the National Library of Medicine, and RLIN are now online.

The last few years have seen internationalization of the system and commercialization, new constituencies well outside of computer science and electrical engineering, regulatory concerns, and security concerns from businesses and out of a concern for our dependence on this as infrastructure. There are questions of pricing and privacy; all of these things are having a significant impact on the technology evolution plan, and with many different stakeholders there are many divergent views of the right way to deal with various problems. These views have to be heard and compromises worked out.

The recent rash of books about the Internet is indicative of the emerging recognition of this system as a very critical international infrastructure, and not just for the research and education community.

I was astonished to see the CCITT bring up an Internet node; the U.N. has just brought up a node, un.org; IEEE and ACM are bringing their systems up. We are well beyond critical mass now. The 1990s will continue this exponential growth phase. The other scary thing is that we are beginning to see experimentation with packet voice and packet video. I fully anticipate that an Internet TV guide will show up in the next couple of years.

I think this kind of phenomenon is going to exacerbate the need for understanding the economics of these systems and how to deal with charging for use of resources. I hesitate to speculate; currently where charges are made they are a fixed price based on the size of the access pipe. It is possible that the continuous transmission requirements of sound and video will require different charging because you are not getting statistical sharing during continuous broadcasting. In the case of multicasting, one packet is multiplied many times. Things like this weren't contemplated when the flat-rate charging algorithms were developed, so the service providers may have to reexamine their charging policies.

Concurrent with the exponential explosion in Internet use has come the recognition that there is a real community out there. The community now needs to recognize that it exists, that it has a diversity of interests, and that it has responsibilities to those who are dependent on the continued health of the network. The Internet Society was founded in January 1992. With assistance from the Federal Networking Council, the Internet Society supports the IETF and IAB and educates the broad community by holding conferences and workshops, by proselytizing, and by making information available.

I had certain technical ambitions when this project started, but they were all oriented toward highly flexible, dynamic communication for military application, insensitive to differences in technology below the level of the routers. I have been extremely pleased with the robustness of the system and its ability to adapt to new communications technology.

One of the main goals of the project was "IP on everything." Whether it is frame relay, ATM, or ISDN, it should always be possible to bring an Internet Protocol up on top of it. We've always been able to get IP to run, so the Internet has satisfied my design criteria. But I didn't have a clue that we would end up with anything like the scale of what we have now, let alone the scale that it's likely to reach by the end of the decade.

On scaling

The somewhat embarrassing thing is that the network address space is under pressure now. The original design of 1973 and 1974 contemplated a total of 256 networks. There was only one LAN at PARC, and all the other networks were regional or nationwide networks. We didn't think there would be more than 256 research networks involved. When it became clear there would be a lot of local area networks, we invented the concept of Class A, B, and C addresses. In Class C there were several million network IDs. But the problem that was not foreseen was that the routing protocols and Internet topology were not well suited for handling an extremely large number of network IDs. So people preferred to use Class B and subnetting instead. We have a rather sparsely allocated address space in the current Internet design, with Class B allocated to excess and Class A and C allocated only lightly.

The lesson is that there is a complex interaction between routing protocols, topology, and scaling, and that determines what Internet routing structure will be necessary for the next ten to twenty years.

When I was chairman of the Internet Activities Board and went to the IETF and IAB to characterize the problem, it was clear that the solution had to be incrementally deployable. You can deploy something in parallel, but then how do the new and old interwork? We are seeing proposals of varying kinds to deal with the problem. Some kind of backward compatibility is highly desirable until you can't assign 32-bit address space. Translating gateways have the defect that when you're halfway through, half the community is transitioned and half isn't, and all the traffic between the two has to go through the translating gateway and it's hard to have enough resources to do this.

It's still a little early to tell how well the alternatives will satisfy the requirements. We are also dealing not only with the scaling problem, but also with the need not to foreclose important new features, such as concepts of flows, the ability to handle multicasting, and concepts of accounting.

I think that as a community we sense varying degrees of pressure for a workable set of solutions. The people who will be most instrumental in this transition will be the vendors of routing equipment and host software, and the offerers of Internet services. It's the people who offer Internet services who have the greatest stake in assuring that Internet operation continues without loss of connectivity, since the value of their service is a function of how many places you can communicate with. The deployability of alternative solutions will determine which is the most attractive. So the transition process is very important.

On use by other networks

The Domain Name System (DNS) has been a key to the scaling of the Internet, allowing it to include non-Internet email systems and solving the problem of name-to-address mapping in a smooth scalable way. Paul Mockapetris deserves enormous credit for the elegant design of the DNS, on which we are still very dependent. Its primary goal was to solve the problems with the host.txt files and to get rid of centralized management. Support for Mail eXchange (MX) was added after the fact, in a second phase.

Once you get a sufficient degree of connectivity, it becomes more advantageous to link to this highly connected thing and tunnel through it rather than to build a system in parallel. So BITNET, FidoNet, AppleTalk, SNA, Novell IPX, and DECNet tunneling are a consequence of the enormous connectivity of the Internet.

The Internet has become a test bed for development of other protocols. Since there was no lower level OSI infrastructure available, Marshall Rose proposed that the Internet could be used to try out X.400 and X.500. In RFC 1006, he proposed that we emulate TPO on top of TCP, and so there was a conscious decision to help higher-level OSI protocols to be deployed in live environments before the lower-level protocols were available.

It seems likely that the Internet will continue to be the environment of choice for the deployment of new protocols and for the linking of diverse systems in the academic, government, and business sectors for the remainder of this decade and well into the next.

Copyright (C) 1993 Vinton Cerf. All rights reserved. May be reproduced in any medium for noncommercial purposes.