

Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn,
Leonard Kleinrock, Daniel C. Lynch, Jon Postel,
Lawrence G. Roberts, Stephen S. Wolff

The Past and Future History of the **INTERNET**

the science of future technology

THE INTERNET HAS REVOLUTIONIZED THE COMPUTER AND COMMUNICATIONS world like nothing before. The telegraph, telephone, radio, and computer have all set the stage for the Internet's unprecedented integration of capabilities. The Internet is at once a worldwide broadcasting capability, a mechanism for information dissemination, and a medium for collaboration and interaction between individuals and their computers without regard for geographic location.

The Internet also represents one of the most successful examples of sustained investment and commitment to research and development in information infrastructure. Beginning with early research in packet switching, the government, industry, and academia have been partners in evolving and deploying this exciting new technology. Today, terms like "leiner@mcc.com" and "http://www.acm.org" trip lightly off the tongue of random people on the street.¹

The Internet today is a widespread information infrastructure, the initial prototype of what is often called the National (or Global or Galactic) Information Infrastructure. Its history is complex and involves

many aspects—technological, organizational, and community. And its influence reaches not only to the technical fields of computer communications but throughout society as we move toward increasing use of online tools to accomplish electronic commerce, information acquisition, and community operations.²

Origins

The first recorded description of the social interactions that could be enabled through networking was a series of memos written August 1962 by J.C.R. Licklider of MIT, discussing his "Galactic Network" concept [6]. Licklider envisioned a globally interconnected set of

¹Perhaps this is an exaggeration due to the lead author's residence in Silicon Valley.

²For a more detailed version of this article, see <http://www.isoc.org/internet-history>.

computers through which everyone could quickly access data and programs from any site. In spirit, the concept was much like the Internet today. While at DARPA,³ he convinced the people who would be his successors there—Ivan Sutherland, Bob Taylor, and MIT researcher Lawrence G. Roberts—of the importance of this networking concept.

Leonard Kleinrock of MIT published the first paper on packet switching theory in July 1961 [5]. Kleinrock convinced Roberts of the theoretical feasibility of communications using packets rather than circuits—a major step toward computer networking. The other key step was to make the computers talk to each other. Exploring this idea in 1965 while working with Thomas Merrill, Roberts connected the TX-2 computer in Massachusetts to the Q-32 in California through a low-speed dial-up telephone line [8], creating the first-ever (though small) wide-area computer network. The result of this experiment: confirmation that time-sharing computers could work well together, running programs and retrieving data as necessary on remote machines, but that the circuit-switched telephone system was totally inadequate for the job. Thus confirmed was Kleinrock's conviction of the need for packet switching.

In late 1966, Roberts went to DARPA to develop the computer network concept and quickly put together a plan for the ARPANET, publishing it in 1967 [7]. Bolt, Beranek and Newman Corp. (BBN), under Frank Heart's leadership, developed the ARPANET switches (called IMPs), with Robert Kahn responsible for overall system design. Howard Frank and his team at Network Analysis Corp. worked with Roberts to optimize the network topology and economics. Due to Kleinrock's early development of packet switching theory and his focus on analysis, design, and measurement, his Network Measurement Center at UCLA was selected as the first node on the ARPANET. All this came together September 1969 when BBN installed the first switch at UCLA and **the first host computer was connected**. In December 1970, the Network Working Group (NWG) under Steve Crocker finished the initial ARPANET host-to-host protocol, called the Network Control Protocol (NCP). As the ARPANET sites completed implementing NCP during 1971–1972, network users finally could begin to develop applications.

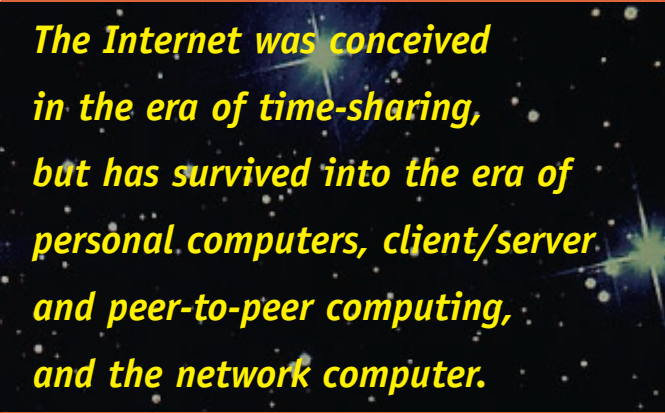
In October 1972, a large, successful demonstration of the ARPANET took place—the first public demon-

stration of this new network technology. Also in 1972, electronic mail, the initial “hot” application, was introduced. In March, Ray Tomlinson of BBN wrote the basic email message send-and-read software, motivated by ARPANET developers' need for an easy coordination mechanism. From there, email took off as the most popular network application and as a harbinger of the kind of people-to-people communication activity we see on the World-Wide Web today.

Initial Internetting Concepts

The original ARPANET grew into the Internet based on the idea that there would be multiple independent networks of rather arbitrary design. Beginning with the ARPANET as the pioneering packet-switching network, it soon grew to include packet satellite networks, ground-based packet radio networks, and other networks. Today's Internet embodies a key underlying technical idea: open-architecture networking. In this approach, the choice of any individual network technology is not dictated by a particular network architecture but can be selected freely by a provider and made to interwork with the other networks through a meta-level “internetworking architecture.” Each network can be designed to fit a specific environment and user requirements.

The idea of open-architecture networking—introduced by Kahn in late 1972 shortly after arriving at DARPA—was guided by four critical ground rules:



*The Internet was conceived
in the era of time-sharing,
but has survived into the era of
personal computers, client/server
and peer-to-peer computing,
and the network computer.*

- Each distinct network had to stand on its own, and no internal changes could be required of any such network before being connected to the Internet.
- Communications would be on a best-effort basis. If a packet didn't make it to the final destination, it would quickly be retransmitted from the source.
- Black boxes (later called gateways and routers) would be used to connect the networks. No infor-

³The Advanced Research Projects Agency (ARPA) changed its name to Defense Advanced Research Projects Agency (DARPA) in 1971, back to ARPA in 1993, and back to DARPA in 1996. We refer throughout to DARPA, the current name.

mation would be retained by the gateways about individual flows of packets passing through them, keeping them simple and avoiding complicated adaptation and recovery from various failure modes.

- There would be no global control at the operations level.

Kahn first began work on a communications-oriented set of operating system principles while at BBN [4]. After joining DARPA and initiating the Internet program, he asked Vinton Cerf (then at Stanford University) to work with him on the detailed design of the protocol. Cerf had been deeply involved in the original NCP design and development and was already knowledgeable about interfacing to existing operating systems. So, armed with Kahn's architectural approach to communications and with Cerf's NCP experience, these two teamed up to spell out the details of what became the Transmission Control Protocol/Internet

The most pressing question for the future of the Internet is not how the technology will change, but how the process of change and evolution itself will be managed.

Protocol (TCP/IP).

The original Cerf/Kahn paper [1] on the Internet described a protocol, called TCP, providing all the Internet's transport and forwarding services. Kahn had intended that TCP would support a range of transport services—from the totally reliable sequenced delivery of data (virtual circuit model) to a datagram service in which the application made direct use of the underlying network service, a process that could imply occasional lost, corrupted, or reordered packets.

However, the initial effort to implement TCP resulted in a version allowing only virtual circuits. This model worked fine for file transfer and remote login applications, but some of the early work on advanced network applications, particularly packet voice in the 1970s, made clear that in some cases packet losses should not be corrected by TCP but left to the application to deal with. This insight led to a reorganization of the original TCP into two protocols—the simple IP providing only for addressing and

forwarding of individual packets and the separate TCP concerned with such service features as flow control and recovery from lost packets. For applications that did not want the services of TCP, an alternative called the User Datagram Protocol (UDP) was added to provide direct access to the basic IP service.

In addition to email, file transfer, and remote login, other applications were proposed in the early days of the Internet, including packet-based voice communication (the precursor of Internet telephony), various models of file and disk sharing, and early “worm” programs illustrating the concept of agents (and viruses). The Internet was not designed for just one application but as a general infrastructure on which new applications could be conceived, exemplified later by the emergence of the Web. The general-purpose nature of the service provided by TCP and IP makes this possible.

Proving the Ideas

DARPA funded three efforts to implement TCP: Stanford (Cerf), BBN (Tomlinson), and University College London (Peter Kirstein). The Stanford team produced a detailed specification, yielding within about a year three independent interoperable implementations of TCP.

This was the beginning of long-term experimentation and development of Internet concepts and technology—along with the constituent networking technologies [3]. Each expansion has involved new challenges. For example, the early implementations of TCP were done for large time-sharing systems. When desktop computers first appeared, it was thought by some that TCP was too big and complex to run on a personal computer. But David Clark and his research group at MIT produced an implementation first for the Xerox Alto (the early personal workstation developed at Xerox PARC) and then for the IBM PC, showing that workstations, as well as large time-sharing systems, could be part of the Internet.

Widespread development of local-area networks (LANs), PCs, and workstations in the 1980s allowed the nascent Internet to flourish. Ethernet technology (developed by Bob Metcalfe at Xerox PARC in 1973) is now probably the dominant network technology in the Internet and PCs and workstations the dominant computers. The increasing scale of the Internet also resulted in several new approaches. For example, the Domain Name System was invented (by Paul Mockapetris, then at USC's Information Sciences Institute) to provide a scalable mechanism for resolving hierarchical host names (e.g., www.acm.org) into Internet addresses. The requirement for scalable routing

approaches led to a hierarchical model of routing, with an Interior Gateway Protocol (IGP) used inside each region of the Internet and an Exterior Gateway Protocol (EGP) used to tie the regions together. New approaches for address aggregation, particularly classless interdomain routing (CIDR), were recently introduced to control the size of router tables.

Another major challenge was how to propagate the changes to the software, particularly host software. DARPA supported the University of California at Berkeley to investigate modifications to the Unix operating system, including incorporating TCP/IP developed at BBN. Although Berkeley later rewrote the BBN code to more efficiently fit into the Unix system and kernel, incorporation of TCP/IP into the Unix BSD system proved a critical element in dispersing the protocols to the research community. Much of the computer science community began using Unix BSD in their day-to-day computing environments. Looking back, the strategy of incorporating Internet protocols into a supported operating system for the research community was a key element in the Internet's successful widespread adoption.

TCP/IP was adopted as a defense standard in 1980, enabling the defense community to begin sharing the DARPA Internet technology base and leading directly to the partitioning of the military and non-military communities. By 1983, ARPANET was being used by a significant number of defense R&D and operational organizations. The transition of ARPANET from NCP to TCP/IP in 1983 permitted it to be split into a MILNET supporting operational requirements and an ARPANET supporting research needs.

Thus, by 1985, the Internet was established as a technology supporting a broad community of researchers and developers and was beginning to be used by other communities for daily computer communications. Email was being used broadly across several communities, often with different systems, demonstrating the utility of broad-based electronic communications between people.

Transition to Widespread Infrastructure

At the same time Internet technology was being experimentally validated and widely used by a subset of computer science researchers, other networks and networking technologies were being pursued. The usefulness of computer networking—especially email—demonstrated by DARPA and Department of Defense contractors on the ARPANET was not lost on other communities and disciplines, so that by the mid-

1970s computer networks began springing up wherever funding was found for the purpose, including the Department of Energy's MFENET and HEPNET, NASA's SPAN, the computer science community's CSNET, the academic community's BITNET, and USENET based on Unix UUCP protocols. Commercial networking technologies were being pursued as well, including IBM's SNA, Xerox's XNS, and Digital Equipment Corp.'s DECNET.

It remained for the British JANET (1984) and U.S. NSFNET (1985) programs to explicitly announce their intent to serve the entire higher education community, regardless of discipline. In 1985, Dennis Jennings came from Ireland for a year to lead the National Science Foundation's NSFNET program. He helped NSF make a critical decision—that TCP/IP would be mandatory for NSFNET. And when Stephen Wolff took over the NSFNET program in 1986, he recognized the need for a wide-area networking infrastructure to support the general academic and research community, as well as the need to develop a strategy for establishing such infrastructure to ultimately be independent of direct federal funding.

Policies and strategies were adopted to achieve that end. So while federal agencies shared the cost of common infrastructure, such as trans-oceanic circuits, NSF encouraged regional networks of the NSFNET to seek commercial, non-academic customers. And NSF enforced an acceptable-use policy, prohibiting Backbone use for purposes "not in support of research and education." The predictable (and intended) result of encouraging commercial network traffic at the local and regional levels, while denying its access to national-scale transport, was the emergence and growth of "private," competitive, long-haul networks, such as PSI, UUNET, ANS CO+RE, and (later) others.

NSF's privatization policy culminated in April 1995 with the defunding of the NSFNET Backbone. The funds thereby recovered were (competitively) redistributed to regional networks to buy national-scale Internet connectivity from the now numerous, private, long-haul networks. The Backbone had made the transition from a network built from routers out of the research community (David Mills's "Fuzzball" routers) to commercial equipment. In its eight-and-a-half-year lifespan, the Backbone had grown from six nodes with 56Kbps links to 21 nodes with multiple 45Mbps. It also saw the Internet grow to more than 50,000 networks on all seven continents and outer space (with 29,000 networks in the U.S.).

Such was the weight of the NSFNET program's

ecumenism and funding (\$200 million, 1986–1995) and the quality of the protocols themselves that by 1990 when the ARPANET itself was finally decommissioned, TCP/IP had supplanted or marginalized most other wide-area computer network protocols worldwide, and IP was on its way to becoming *the* bearer service for the Global Information Infrastructure.

Documentation's Key Role

A key to the rapid growth of the Internet has been free and open access to the basic documents, especially the specifications of the protocols. The beginnings of the ARPANET and the Internet in the university research community promoted the academic tradition of open publication of ideas and results. However, the normal cycle of traditional academic publication was too formal and too slow for the dynamic exchange of ideas essential to creating networks. In 1969, a key step was taken by S. Crocker (then at UCLA) in establishing the request for comments (or RFC) series of notes [2]. These memos were intended to be an informal fast means of distribution for sharing ideas among network researchers. At first the RFCs were printed on paper and distributed via postal mail. As the File Transfer Protocol (FTP) came into use, the RFCs were prepared as online files and accessed via FTP. Now, the RFCs are easily accessed via the Web at dozens of sites around the world. SRI, in its role as Network Information Center, maintained the online directories. Jon Postel acted as RFC editor and as manager of centralized administration of required protocol number assignments—roles he continues to this day.

The effect of the RFCs was to create a positive feedback loop, so ideas or proposals presented in one RFC would trigger other RFCs. When consensus (or at least a consistent set of ideas) would come together, a specification document would be prepared; such specifications would then be used as the basis for implementations by the various research teams. The RFCs are now viewed as the “documents of record” in the Internet engineering and standards community and will continue to be critical to future net evolution while furthering the net’s initial role of sharing information about its own design and operations.

Formation of a Broad Community

The Internet is as much a collection of communities as a collection of technologies, and its success is largely attributable to satisfying basic community needs as well as utilizing the community effectively to push the

infrastructure forward. Community spirit has a long history, beginning with the early ARPANET, whose early researchers worked as a close-knit community (the ARPANET Working Group) to accomplish the initial demonstrations of packet switching technology. Likewise, the packet satellite, packet radio, and other DARPA computer science research programs were multi-contractor collaborative activities that used any available mechanisms to coordinate their efforts, starting with email, and adding file sharing, remote access, and eventually Web capabilities.

In the late 1970s, recognizing that the growth of the Internet was accompanied by the growth of the interested research community and therefore an increased need for coordination mechanisms, Cerf, then manager of the DARPA Internet program, formed several coordination bodies, including the Internet Configuration Control Board (ICCB), chaired by Clark. The ICCB was an invitation-only body assisting Cerf in managing the burgeoning Internet activity.

In 1983, when Barry Leiner took over management of the Internet program at DARPA, he and Clark recognized that the continuing growth of the Internet community demanded a restructuring of the coordination mechanisms.

The ICCB was disbanded and replaced by a structure of Task Forces, each focused on a particular area of the technology (e.g., routers and end-to-end protocols). The Internet Activities Board (IAB) included the chairs of the Task Forces.

After some changing membership on the IAB, Phill Gross became chair of a revitalized Internet Engineering Task Force (IETF)—at the time only one of the IAB Task Forces. The growth of the Internet in the mid-1980s resulted in vastly increased attendance at IETF meetings, and Gross had to create substructure to the IETF in the form of working groups.

The expanded community also meant that DARPA was no longer the only major player when it came to funding the Internet. In addition to NSFNET and the various U.S. and international government-funded activities, interest in the commercial sector was beginning to grow. And in 1985, when both Kahn and Leiner left DARPA, there was a significant decrease in DARPA's Internet activity. The IAB was left without a primary sponsor and so increasingly assumed the mantle of leadership.

Continued growth resulted in even further substructure within both the IAB and IETF, while growth in the commercial sector brought increased concern

“
In the future, computers will shop for us. You will log into a virtual supermarket and order food, they will send it to you the next day. School will change too. You could have school at home and fax your homework in, but you won't make any friends that way. We will have to make friends on the Internet. Libraries will still be there, but not many people will visit them anymore and they will be knocked down.”

—NICHOLAS PHIBBS, AGE 12
SURREY, UK

regarding the standards process. The twin motivations of making the process open and fair and the need to win Internet community support eventually led in 1991 to formation of the Internet Society, under the auspices of Kahn's Corporation for National Research Initiatives (CNRI) and the leadership of Cerf, who was then with CNRI.

In 1992, yet another reorganization took place. The IAB was reorganized and renamed the Internet Architecture Board. A more peer-like relationship was defined between the new IAB and Internet Engineering Steering Group (IESG), with the IETF and IESG taking greater responsibility for approving standards. Ultimately, a cooperative and mutually supportive relationship was formed among the IAB, IETF, and the Internet Society.

The Web's recent development and widespread deployment brings a new community, as many of the people now working on the Web didn't view themselves primarily as network researchers and developers. Therefore, in 1995, a new coordination organization was formed—the World-Wide Web Consortium (W3C), initially led from MIT's Laboratory for Computer Science by Al Vezza and Tim Berners-Lee, the Web's inventor. Today, the W3C is responsible for evolving the various protocols and standards associated with the Web.

Commercialization

Commercialization of the Internet has involved not only development of competitive, private network services, but commercial products implementing Internet

technology. In the early 1980s, dozens of vendors were incorporating TCP/IP into their products because they saw buyers for that approach to networking. Unfortunately, they lacked real information about how the technology was supposed to work and how their customers planned to use the approach.

In 1985, recognizing the lack of available information and appropriate training, Daniel Lynch in cooperation with the IAB arranged a three-day workshop for all vendors to learn how TCP/IP worked and what it still could not do well. Speakers were mostly from the DARPA research community where they had developed these protocols and used them in day-to-day work. Approximately 250 vendor representatives heard 50 inventors and experimenters.

The first Interop trade show, September 1988, demonstrated interoperability between vendor products and was attended by 50 companies and 5,000 engineers from potential customer organizations. Interop has grown immensely since then, and today is an annual event in seven locations around the world for an audience of more than 250,000 who want to learn which products seamlessly work with which other products and about the latest technology.

In the last few years, we have seen a new phase of commercialization. Originally, commercial efforts mainly comprised vendors providing the basic networking products and service providers offering connectivity and basic Internet services. The Internet has now become almost a “commodity” service, and much of the latest attention has been on the use of this global information infrastructure as support

for other commercial services.

This activity has been accelerated by the widespread and rapid adoption of browsers and Web technology, giving users easy access to information linked around the globe. Products are available for finding, sending, and retrieving that information, and many of the latest developments seek to provide increasingly sophisticated information services on top of basic Internet data communications.

History of the Future

The Internet was conceived in the era of time-sharing, but has survived into the era of personal computers, client/server and peer-to-peer computing, and the network computer. It was designed before LANs existed, but has evolved to accommodate LANs as well as more recent ATM and frame-switched services. It was envisioned as supporting a range of functions—from file sharing and remote login to resource sharing and collaboration, and has spawned email and more recently the Web. But most important, it started as the creation of a small band of dedicated researchers and has grown to be a commercial success with billions of dollars invested annually.

One should not conclude that the Internet is complete. The Internet is a creature of the computer, not the traditional networks of the telephone or television industries. It will—indeed it must—continue changing at the speed of the computer industry to remain relevant. It is now changing to provide such new services as real-time transport, supporting, for example, audio and video streams. The availability of pervasive networking—that is, the Internet itself—along with powerful affordable computing and communications in portable form (e.g., laptop computers, two-way pagers, PDAs, cellular phones) makes possible a new paradigm of nomadic computing and communications.

This evolution will bring us new applications—Internet telephone and, further out, Internet television. It will also permit more sophisticated forms of pricing and cost recovery, a perhaps painful requirement in this commercial world. It is changing to accommodate yet another generation of underlying network technologies with different characteristics and requirements—from broadband residential access to satellites. New modes of access and new forms of service will spawn new applications that in turn will drive further evolution of the net itself.

The most pressing question for the future of the Internet is not how the technology will change, but how the process of change and evolution itself will be

managed. Internet architecture has always been driven by a core group of designers, but the form of that group has changed as the number of interested outside parties has grown. With the success of the Internet has come a proliferation of stakeholders—now with an economic as well as an intellectual investment in the network. We see, for example, in the debates over control of the domain namespace and the form of the next-generation IP addresses a struggle to find the next social structure to guide the Internet. However, that structure is more difficult to define, given the large number of stakeholders. The industry also struggles to find the economic rationale for the huge investment needed for future growth to, for example, upgrade residential access to more suitable technology. If the Internet stumbles, it will not be because we lack technology, vision, or motivation but because we cannot set a direction and march collectively into the future. ■

References

1. Cerf, V.G., and Kahn, R.E. A protocol for packet network interconnection. *IEEE Trans. Comm. Tech* 5 (May 1974), 627–641.
2. Crocker, S. Host software. RFC 001. Apr. 7, 1969.
3. Kahn, R. guest ed., Uncapher, K., van Trees, H., assoc. guest eds. Special Issue on Packet Communication Networks. *Proceedings of the IEEE* 66, 11 (Nov. 1978).
4. Kahn, R. Communications Principles for Operating Systems. Internal BBN memorandum, Jan. 1972.
5. Kleinrock, L. Information Flow in Large Communication Nets. RLE Quarterly Progress Report, July 1961.
6. Licklider, J.C.R., and Clark, W. On-Line Man-Computer Communication. Aug. 1962.
7. Roberts, L. Multiple Computer Networks and Intercomputer Communication. In *Proceedings of the ACM Gatlinburg Conference* (Oct. 1967).
8. Roberts, L., and Merrill, T. Toward a Cooperative Network of Time-Shared Computers. In *Proceedings of the Fall AFIPS Conference* (Oct. 1966).

BARRY M. LEINER (leiner@mcc.com) is Vice President of Microelectronics and Computer Technology Corp.

VINTON G. CERF (vcerf@mci.net) is Senior Vice President, Internet Architecture and Engineering, at MCI Communications Corp.

DAVID D. CLARK (ddc@lcs.mit.edu) is a senior research scientist at the MIT Laboratory for Computer Science.

ROBERT E. KAHN (rkahn@cnri.reston.va.us) is President of the Corporation for National Research Initiatives.

LEONARD KLEINROCK (lk@cs.ucla.edu) is a professor of computer science at the University of California, Los Angeles.

DANIEL C. LYNCH (dlynch@cybercash.com) is Chairman of CyberCash Inc. and founder of the Interop networking trade show and conferences.

JON POSTEL (postel@isi.edu) is the Director of the Computer Networks Division of the Information Sciences Institute of the University of Southern California.

LAWRENCE G. ROBERTS (lroberts@atmsys.com) is the President of ATM Systems Division of Connectware Inc.

STEPHEN S. WOLFF (swolff@cisco.com) is with Cisco Systems, Inc.

Copyright held by the authors