

Brief History of the Internet

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

The Internet is at once a world-wide broadcasting capability, a mechanism for information dissemination, and a medium for collaboration and interaction between individuals and their computers without regard for geographic location.

Introduction

The Internet has revolutionized the computer and communications world like nothing before. The invention of the telegraph, telephone, radio, and computer set the stage for this unprecedented integration of capabilities. The Internet is at once a world-wide 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 represents one of the most successful examples of the benefits of sustained investment and commitment to research and development of information infrastructure. Beginning with the early research in packet switching, the government, industry and academia have been partners in evolving and deploying this exciting new technology. Today, terms like "bleiner@computer.org" and "http://www.acm.org" trip lightly off the tongue of the random person on the street.¹

This is intended to be a brief, necessarily cursory and incomplete history. Much material currently exists about the Internet, covering history, technology, and usage. A trip to almost any bookstore will find shelves of material written about the Internet.²

In this paper,³ several of us involved in the development and evolution of the Internet share our views of its origins and history. This history revolves around four distinct aspects. There is the technological evolution that began with early research on packet switching and the ARPANET (and related technologies), and where current research continues to expand the horizons of the infrastructure along several dimensions, such as scale, performance, and higher-level functionality. There is the operations and management aspect of a global and complex operational infrastructure. There is the social aspect, which resulted in a broad community of Internauts working together to create and evolve the technology. And there is the commercialization aspect, resulting in an extremely effective transition of research results into a broadly deployed and available information infrastructure.

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.

Internet 20 YEARS

Origins of the Internet

The first recorded description of the social interactions that could be enabled through networking was a series of memos written by J.C.R. Licklider of MIT in August 1962 discussing his "Galactic Network" concept. He envisioned a globally interconnected set of computers through which everyone could quickly access data and programs from any site. In spirit, the concept was very much like the Internet of today. Licklider was the first head of the computer research program at DARPA,⁴ starting in October 1962. While at DARPA he convinced his successors at DARPA, Ivan Sutherland, Bob Taylor, and MIT researcher Lawrence G. Roberts, of the importance of this networking concept.

Leonard Kleinrock at MIT published the first paper on packet switching theory in July 1961 and the first book on the subject in 1964. Kleinrock convinced Roberts of the theoretical feasibility of communications using packets rather than circuits, which was a major step along the path towards computer networking. The other key step was to make the computers talk together. To explore this, in 1965 working with Thomas Merrill, Roberts connected the TX-2 computer in Mass. to the Q-32 in California with a low speed dial-up telephone line creating the first (however small) wide-area computer network ever built. The result of this experiment was the realization that the time-shared computers could work well together, running programs and retrieving data as necessary on the remote machine, but that the circuit switched telephone system was totally inadequate for the job. Kleinrock's conviction of the need for packet switching was confirmed.

In late 1966 Roberts went to DARPA to develop the computer network concept and quickly put together his plan for the "ARPANET", publishing it in 1967. At the conference where he presented the paper, there was also a paper on a packet network concept from the UK by Donald Davies and Roger Scantlebury of NPL. Scantlebury told Roberts about the NPL work as well as that of Paul Baran and others at RAND. The RAND group had written a paper on packet switching networks for secure voice in the military in 1964. It happened that the work at MIT (1961-1967), at RAND (1962-1965), and at NPL (1964-1967) had all proceeded in parallel without any of the researchers knowing about the other work. The word "packet" was adopted from the work at NPL and the proposed line speed to be used in the ARPANET design was upgraded from 2.4 kbps to 50 kbps.⁵

In August 1968, after Roberts and the DARPA funded community had refined the overall structure and specifications for the ARPANET, an RFQ was released by DARPA for the development of one of the key components, the packet switches called Interface Message Processors (IMP's). The RFQ was won in December 1968 by a group headed by Frank Heart at Bolt Beranek and Newman (BBN). As the BBN team worked on the IMP's with Bob Kahn playing a major role in the overall ARPANET architectural design, the network topology and economics were designed and optimized by Roberts working with Howard Frank and his team at Network Analysis Corporation, and the network measurement system was prepared by Kleinrock's team at UCLA.⁶

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 to be the first node on the ARPANET. All this came together in September 1969 when BBN installed the first IMP at UCLA and the first host computer was connected. Doug Engelbart's project on "Augmentation of Human Intellect" (which included NLS, an early hypertext system) at Stanford Research Institute (SRI) provided a second node. SRI supported the Network Information Center, led by Elizabeth (Jake) Feinler and including functions such as maintaining tables of host name to address mapping as well as a directory of the RFC's.

One month later, when SRI was connected to the ARPANET, the first host-to-host message was sent from Kleinrock's laboratory to SRI. Two more nodes were added at UC Santa Barbara and University of Utah. These last two nodes incorporated application visualization projects, with Glen Culler and Burton Fried at UCSB investigating methods for display of mathematical functions using storage displays to deal with the problem of refresh over the net, and Robert Taylor and Ivan Sutherland at Utah investigating methods of 3-D representations over the net. Thus, by the end of 1969, four host computers were connected together into the initial ARPANET, and the budding Internet was off the ground. Even at this early stage, it should be noted that the networking research incorporated both work on the underlying network and work on how to utilize the network. This tradition continues to this day.

Computers were added quickly to the ARPANET during the following years, and work proceeded on completing a functionally complete Host-to-Host protocol and other network software. In December 1970 the Network Working Group (NWG) working under S. Crocker finished the initial ARPANET Host-to-Host protocol, called the Network Control Protocol (NCP). As the ARPANET sites completed implementing NCP during the period 1971-1972, the network users finally could begin to develop applications.

In October 1972, Kahn organized a large, very successful demonstration of the ARPANET at the International Computer Communication Conference (ICCC). This was the first public demonstration of this new network technology to the public. It was also in 1972 that the initial "hot" application, electronic mail, was introduced. In March Ray Tomlinson at BBN wrote the basic email message send and read software, motivated by the need of the ARPANET developers for an easy coordination mechanism. In July, Roberts expanded its utility by writing the first email utility program to list, selectively read, file, forward, and respond to messages. From there email took off as the largest network application for over a decade. This was a harbinger of the kind of activity we see on the World Wide Web today, namely, the enormous growth of all kinds of "people-to-people" traffic.

The Initial Internetworking Concepts

The original ARPANET grew into the Internet. Internet was 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, but soon to include packet satellite networks, ground-based packet radio networks and other networks. The Internet as we now know it embodies a key underlying technical idea, namely that of open architecture networking. In this approach, the choice of any individual network technology was not dictated by a particular network architecture but rather could be selected freely by a provider and made to interwork with the other networks through a meta-level "Internetworking Architecture". Up until that time there was only one general method for federating networks. This was the traditional circuit switching method where networks would interconnect at the circuit level, passing individual bits on a synchronous basis along a portion of an end-to-end circuit between a pair of end locations. Recall that Kleinrock had shown in 1961 that packet switching was a more efficient switching method. Along with packet switching, special purpose interconnection arrangements between networks were another possibility. While there were other limited ways to interconnect different networks, they required that one be used as a component of the other, rather than acting as a peer of the other in offering end-to-end service.

In an open-architecture network, the individual networks may be separately designed and developed and each may have its own unique interface which it may offer to users and/or other providers. including other Internet providers. Each network can be designed in accordance with

the specific environment and user requirements of that network. There are generally no constraints on the types of network that can be included or on their geographic scope, although certain pragmatic considerations will dictate what makes sense to offer.

The idea of open-architecture networking was first introduced by Kahn shortly after having arrived at DARPA in 1972. This work was originally part of the packet radio program, but subsequently became a separate program in its own right. At the time, the program was called "Interneting". Key to making the packet radio system work was a reliable end-end protocol that could maintain effective communication in the face of jamming and other radio interference, or withstand intermittent blackout such as caused by being in a tunnel or blocked by the local terrain. Kahn first contemplated developing a protocol local only to the packet radio network, since that would avoid having to deal with the multitude of different operating systems, and continuing to use NCP.

However, NCP did not have the ability to address networks (and machines) further downstream than a destination IMP on the ARPANET and thus some change to NCP would also be required. (The assumption was that the ARPANET was not changeable in this regard). NCP relied on ARPANET to provide end-to-end reliability. If any packets were lost, the protocol (and presumably any applications it supported) would come to a grinding halt. In this model NCP had no end-end host error control, since the ARPANET was to be the only network in existence and it would be so reliable that no error control would be required on the part of the hosts. Thus, Kahn decided to develop a new version of the protocol which could meet the needs of an open-architecture network environment. This protocol would eventually be called the Transmission Control Protocol/Internet Protocol (TCP/IP). While NCP tended to act like a device driver, the new protocol would be more like a communications protocol.

Four ground rules were critical to Kahn's early thinking:

- Each distinct network would have to stand on its own and no internal changes could be required to any such network to connect it to the Internet.
- Communications would be on a best effort basis. If a packet didn't make it to the final destination, it would shortly be retransmitted from the source.
- Black boxes would be used to connect the networks; these would later be called gateways and routers. There would be no information retained by the gateways about the individual flows of packets passing through them, thereby keeping them simple and avoiding complicated adaptation and recovery from various failure modes.
- There would be no global control at the operations level.

Other key issues that needed to be addressed were:

- Algorithms to prevent lost packets from permanently disabling communications and enabling them to be successfully retransmitted from the source.
- Providing for host-to-host "pipelining" so that multiple packets could be enroute from source to destination at the discretion of the participating hosts, if the intermediate networks allowed it.
- Gateway functions to allow it to forward packets appropriately. This included interpreting IP headers for routing, handling interfaces, breaking packets into smaller pieces if necessary, etc.
- The need for end-end checksums, reassembly of packets from fragments and detection of duplicates, if any.
- The need for global addressing

- Techniques for host-to-host flow control.
- Interfacing with the various operating systems
- There were also other concerns, such as implementation efficiency, internetwork performance, but these were secondary considerations at first.

Kahn began work on a communications-oriented set of operating system principles while at BBN and documented some of his early thoughts in an internal BBN memorandum entitled "Communications Principles for Operating Systems". At this point he realized it would be necessary to learn the implementation details of each operating system to have a chance to embed any new protocols in an efficient way. Thus, in the spring of 1973, after starting the internetworking effort, he asked Vint Cerf (then at Stanford) to work with him on the detailed design of the protocol. Cerf had been intimately involved in the original NCP design and development and already had the knowledge about interfacing to existing operating systems. So armed with Kahn's architectural approach to the communications side and with Cerf's NCP experience, they teamed up to spell out the details of what became TCP/IP.

The give and take was highly productive and the first written version⁷ of the resulting approach was distributed at a special meeting of the International Network Working Group (INWG) which had been set up at a conference at Sussex University in September 1973. Cerf had been invited to chair this group and used the occasion to hold a meeting of INWG members who were heavily represented at the Sussex Conference.

Some basic approaches emerged from this collaboration between Kahn and Cerf:

- Communication between two processes would logically consist of a very long stream of bytes (they called them octets). The position of any octet in the stream would be used to identify it.
- Flow control would be done by using sliding windows and acknowledgments (acks). The destination could select when to acknowledge and each ack returned would be cumulative for all packets received to that point.
- It was left open as to exactly how the source and destination would agree on the parameters of the windowing to be used. Defaults were used initially.
- Although Ethernet was under development at Xerox PARC at that time, the proliferation of LANs were not envisioned at the time, much less PCs and workstations. The original model was national level networks like ARPANET of which only a relatively small number were expected to exist. Thus a 32 bit IP address was used of which the first⁸ bits signified the network and the remaining 24 bits designated the host on that network. This assumption, that 256 networks would be sufficient for the foreseeable future, was clearly in need of reconsideration when LANs began to appear in the late 1970s.

The original Cerf/Kahn paper on the Internet described one protocol, called TCP, which provided all the transport and forwarding services in the Internet. Kahn had intended that the TCP protocol 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, which might imply occasional lost, corrupted or reordered packets. However, the initial effort to implement TCP resulted in a version that only allowed for virtual circuits. This model worked fine for file transfer and remote login applications, but some of the early work on advanced network applications, in particular packet voice in the 1970s, made clear that in some cases packet losses should not be corrected by TCP, but should be left to the application to deal with. This led to a reorganization of the original TCP into two protocols, the simple IP which provided only for addressing and forwarding of individual packets, and the separate TCP, which was

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.