Subject: A short history of the Internet (Feb 1993) (Bruce Sterling)

From THE MAGAZINE OF FANTASY AND SCIENCE FICTION, Feb 1993.
F&SF, Box 56, Cornwall CT 06753 $26/yr USA $31/yr other

Some thirty years ago, the RAND Corporation, America's foremost Cold War think-tank, faced a strange strategic problem. How could the US authorities successfully communicate after a nuclear war?

Postnuclear America would need a command-and-control network, linked from city to city, state to state, base to base. But no matter how thoroughly that network was armored or protected, its switches and wiring would always be vulnerable to the impact of atomic bombs. A nuclear attack would reduce any conceivable network to tatters.

And how would the network itself be commanded and controlled? Any central authority, any network central citadel, would be an obvious and immediate target for an enemy missile. The center of the network would be the very first place to go.

RAND mulled over this grim puzzle in deep military secrecy, and arrived at a daring solution. The RAND proposal (the brainchild of RAND staffer Paul Baran) was made public in 1964. In the first place, the network would *have no central authority.* Furthermore, it would be *designed from the beginning to operate while in tatters.*

The principles were simple. The network itself would be assumed to be unreliable at all times. It would be designed from the get-go to transcend its own unreliability. All the nodes in the network would be equal in status to all other nodes, each node with its own authority to originate, pass, and receive messages. The messages themselves would be divided into packets, each packet separately addressed. Each packet would begin at some specified source node, and end at some other specified destination node. Each packet would wind its way through the network on an individual basis.

The particular route that the packet took would be unimportant. Only final results would count. Basically, the packet would be tossed like a hot potato from node to node to node, more or less in the direction of its destination, until it ended up in the proper place. If big pieces of the network had been blown away, that simply wouldn't matter; the packets would still stay airborne, lateralled wildly across the field by whatever nodes happened to survive. This rather haphazard delivery system might be "inefficient" in the usual sense (especially compared to, say, the telephone system) -- but it would be extremely rugged.

During the 60s, this intriguing concept of a decentralized, blastproof, packet-switching network was kicked around by RAND, MIT and UCLA. The National Physical Laboratory in Great Britain set up the first test network on these principles in 1968. Shortly afterward, the Pentagon's Advanced Research Projects Agency decided to fund a larger, more ambitious project in the

USA. The nodes of the network were to be high-speed supercomputers (or what passed for supercomputers at the time). These were rare and valuable machines which were in real need of good solid networking, for the sake of national research-and-development projects.

In fall 1969, the first such node was installed in UCLA. By December 1969, there were four nodes on the infant network, which was named ARPANET, after its Pentagon sponsor.

The four computers could transfer data on dedicated high-speed transmission lines. They could even be programmed remotely from the other nodes. Thanks to ARPANET, scientists and researchers could share one another's computer facilities by long-distance. This was a very handy service, for computer-time was precious in the early '70s. In 1971 there were fifteen nodes in ARPANET; by 1972, thirty-seven nodes. And it was good.

By the second year of operation, however, an odd fact became clear. ARPANET's users had warped the computer-sharing network into a dedicated, high-speed, federally subsidized electronic post-office. The main traffic on ARPANET was not long-distance computing. Instead, it was news and personal messages. Researchers were using ARPANET to collaborate on projects, to trade notes on work, and eventually, to downright gossip and schmooze. People had their own personal user accounts on the ARPANET computers, and their own personal addresses for electronic mail. Not only were they using ARPANET for person-to-person communication, but they were very enthusiastic about this particular service -- far more enthusiastic than they were about long-distance computation.

It wasn't long before the invention of the mailing-list, an ARPANET broadcasting technique in which an identical message could be sent automatically to large numbers of network subscribers. Interestingly, one of the first really big mailing-lists was "SF-LOVERS," for science fiction fans. Discussing science fiction on the network was not work-related and was frowned upon by many ARPANET computer administrators, but this didn't stop it from happening.

Throughout the '70s, ARPA's network grew. Its decentralized structure made expansion easy. Unlike standard corporate computer networks, the ARPA network could accommodate many different kinds of machine. As long as individual machines could speak the packet-switching lingua franca of the new, anarchic network, their brand-names, and their content, and even their ownership, were irrelevant.

The ARPA's original standard for communication was known as NCP, "Network Control Protocol," but as time passed and the technique advanced, NCP was superceded by a higher-level, more sophisticated standard known as TCP/IP. TCP, or "Transmission Control Protocol," converts messages into streams of packets at the source, then reassembles them back into messages at the destination. IP, or "Internet Protocol," handles the addressing, seeing to it that packets are routed across multiple nodes and even across multiple networks with multiple standards -- not only ARPA's pioneering NCP standard, but others like Ethernet, FDDI, and X.25.

As early as 1977, TCP/IP was being used by other networks to link to ARPANET. ARPANET itself remained fairly tightly controlled, at least until 1983, when its military segment broke off and became MILNET. But TCP/IP linked them all. And ARPANET itself, though it was growing, became a smaller and smaller neighborhood amid the vastly growing galaxy of other linked machines.

As the '70s and '80s advanced, many very different social groups found themselves in possession of powerful computers. It was fairly easy to link these computers to the growing network-of-networks. As the use of TCP/IP became more common, entire other networks fell into the digital embrace of the Internet, and messily adhered. Since the software called TCP/IP was public-domain, and the basic technology was decentralized and rather anarchic by its very nature, it was difficult to stop people from barging in and linking up somewhere-or-other. In point of fact, nobody *wanted* to stop them from joining this branching complex of networks, which came to be known as the "Internet."

Connecting to the Internet cost the taxpayer little or nothing, since each node was independent, and had to handle its own financing and its own technical requirements. The more, the merrier. Like the phone network, the computer network became steadily more valuable as it embraced larger and larger territories of people and resources.

A fax machine is only valuable if *everybody else* has a fax machine. Until they do, a fax machine is just a curiosity. ARPANET, too, was a curiosity for a while. Then computer-networking became an utter necessity.

In 1984 the National Science Foundation got into the act, through its Office of Advanced Scientific Computing. The new NSFNET set a blistering pace for technical advancement, linking newer, faster, shinier supercomputers, through thicker, faster links, upgraded and expanded, again and again, in 1986, 1988, 1990. And other government agencies leapt in: NASA, the National Institutes of Health, the Department of Energy, each of them maintaining a digital satrapy in the Internet confederation.

The nodes in this growing network-of-networks were divvied up into basic varieties. Foreign computers, and a few American ones, chose to be denoted by their geographical locations. The others were grouped by the six basic Internet "domains": gov, mil, edu, com, org and net. (Graceless abbreviations such as this are a standard feature of the TCP/IP protocols.) Gov, Mil, and Edu denoted governmental, military and educational institutions, which were, of course, the pioneers, since ARPANET had begun as a high-tech research exercise in national security. Com, however, stood for "commercial" institutions, which were soon bursting into the network like rodeo bulls, surrounded by a dust-cloud of eager nonprofit "orgs." (The "net" computers served as gateways between networks.)

ARPANET itself formally expired in 1989, a happy victim of its own overwhelming success. Its users scarcely noticed, for ARPANET's functions not only continued but steadily improved. The

use of TCP/IP standards for computer networking is now global. In 1971, a mere twenty-one years ago, there were only four nodes in the ARPANET network. Today there are tens of thousands of nodes in the Internet, scattered over forty-two countries, with more coming on-line every day. Three million, possibly four million people use this gigantic mother-of-all-computer-networks.

The Internet is especially popular among scientists, and is probably the most important scientific instrument of the late twentieth century. The powerful, sophisticated access that it provides to specialized data and personal communication has sped up the pace of scientific research enormously.

The Internet's pace of growth in the early 1990s is spectacular, almost ferocious. It is spreading faster than cellular phones, faster than fax machines. Last year the Internet was growing at a rate of twenty percent a *month.* The number of "host" machines with direct connection to TCP/IP has been doubling every year since 1988. The Internet is moving out of its original base in military and research institutions, into elementary and high schools, as well as into public libraries and the commercial sector.

Why do people want to be "on the Internet?" One of the main reasons is simple freedom. The Internet is a rare example of a true, modern, functional anarchy. There is no "Internet Inc." There are no official censors, no bosses, no board of directors, no stockholders. In principle, any node can speak as a peer to any other node, as long as it obeys the rules of the TCP/IP protocols, which are strictly technical, not social or political. (There has been some struggle over commercial use of the Internet, but that situation is changing as businesses supply their own links).

The Internet is also a bargain. The Internet as a whole, unlike the phone system, doesn't charge for long-distance service. And unlike most commercial computer networks, it doesn't charge for access time, either. In fact the "Internet" itself, which doesn't even officially exist as an entity, never "charges" for anything. Each group of people accessing the Internet is responsible for their own machine and their own section of line.

The Internet's "anarchy" may seem strange or even unnatural, but it makes a certain deep and basic sense. It's rather like the "anarchy" of the English language. Nobody rents English, and nobody owns English. As an English-speaking person, it's up to you to learn how to speak English properly and make whatever use you please of it (though the government provides certain subsidies to help you learn to read and write a bit). Otherwise, everybody just sort of pitches in, and somehow the thing evolves on its own, and somehow turns out workable. And interesting. Fascinating, even. Though a lot of people earn their living from using and exploiting and teaching English, "English" as an institution is public property, a public good. Much the same goes for the Internet. Would English be improved if the "The English Language, Inc." had a board of directors and a chief executive officer, or a President and a Congress? There'd probably be a lot fewer new words in English, and a lot fewer new ideas.

People on the Internet feel much the same way about their own institution. It's an institution that resists institutionalization. The Internet belongs to everyone and no one.

Still, its various interest groups all have a claim. Business people want the Internet put on a sounder financial footing. Government people want the Internet more fully regulated. Academics want it dedicated exclusively to scholarly research. Military people want it spy-proof and secure. And so on and so on.

All these sources of conflict remain in a stumbling balance today, and the Internet, so far, remains in a thrivingly anarchical condition. Once upon a time, the NSFnet's high-speed, high-capacity lines were known as the "Internet Backbone," and their owners could rather lord it over the rest of the Internet; but today there are "backbones" in Canada, Japan, and Europe, and even privately owned commercial Internet backbones specially created for carrying business traffic. Today, even privately owned desktop computers can become Internet nodes. You can carry one under your arm. Soon, perhaps, on your wrist.

But what does one *do* with the Internet? Four things, basically: mail, discussion groups, long-distance computing, and file transfers.

Internet mail is "e-mail," electronic mail, faster by several orders of magnitude than the US Mail, which is scornfully known by Internet regulars as "snailmail." Internet mail is somewhat like fax. It's electronic text. But you don't have to pay for it (at least not directly), and it's global in scope. E-mail can also send software and certain forms of compressed digital imagery. New forms of mail are in the works.

The discussion groups, or "newsgroups," are a world of their own. This world of news, debate and argument is generally known as "USENET. " USENET is, in point of fact, quite different from the Internet. USENET is rather like an enormous billowing crowd of gossipy, news-hungry people, wandering in and through the Internet on their way to various private backyard barbecues. USENET is not so much a physical network as a set of social conventions. In any case, at the moment there are some 2,500 separate newsgroups on USENET, and their discussions generate about 7 million words of typed commentary every single day. Naturally there is a vast amount of talk about computers on USENET, but the variety of subjects discussed is enormous, and it's growing larger all the time. USENET also distributes various free electronic journals and publications.

Both netnews and e-mail are very widely available, even outside the high-speed core of the Internet itself. News and e-mail are easily available over common phone-lines, from Internet fringe-realms like BITnet, UUCP and Fidonet. The last two Internet services, long-distance computing and file transfer, require what is known as "direct Internet access" -- using TCP/IP.

Long-distance computing was an original inspiration for ARPANET and is still a very useful service, at least for some. Programmers can maintain accounts on distant, powerful computers, run programs there or write their own. Scientists can make use of powerful supercomputers a continent away. Libraries offer their electronic card catalogs for free search. Enormous CD-ROM

catalogs are increasingly available through this service. And there are fantastic amounts of free software available.

File transfers allow Internet users to access remote machines and retrieve programs or text. Many Internet computers -- some two thousand of them, so far -- allow any person to access them anonymously, and to simply copy their public files, free of charge. This is no small deal, since entire books can be transferred through direct Internet access in a matter of minutes. Today, in 1992, there are over a million such public files available to anyone who asks for them (and many more millions of files are available to people with accounts). Internet file-transfers are becoming a new form of publishing, in which the reader simply electronically copies the work on demand, in any quantity he or she wants, for free. New Internet programs, such as "archie," "gopher," and "WAIS," have been developed to catalog and explore these enormous archives of material.

The headless, anarchic, million-limbed Internet is spreading like bread-mold. Any computer of sufficient power is a potential spore for the Internet, and today such computers sell for less than $2,000 and are in the hands of people all over the world. ARPA's network, designed to assure control of a ravaged society after a nuclear holocaust, has been superceded by its mutant child the Internet, which is thoroughly out of control, and spreading exponentially through the post-Cold War electronic global village. The spread of the Internet in the 90s resembles the spread of personal computing in the 1970s, though it is even faster and perhaps more important. More important, perhaps, because it may give those personal computers a means of cheap, easy storage and access that is truly planetary in scale.

The future of the Internet bids fair to be bigger and exponentially faster. Commercialization of the Internet is a very hot topic today, with every manner of wild new commercial information-service promised. The federal government, pleased with an unsought success, is also still very much in the act. NREN, the National Research and Education Network, was approved by the US Congress in fall 1991, as a five-year, $2 billion project to upgrade the Internet "backbone." NREN will be some fifty times faster than the fastest network available today, allowing the electronic transfer of the entire Encyclopedia Britannica in one hot second. Computer networks worldwide will feature 3-D animated graphics, radio and cellular phone-links to portable computers, as well as fax, voice, and high-definition television. A multimedia global circus!

Or so it's hoped -- and planned. The real Internet of the future may bear very little resemblance to today's plans. Planning has never seemed to have much to do with the seething, fungal development of the Internet. After all, today's Internet bears little resemblance to those original grim plans for RAND's post-holocaust command grid. It's a fine and happy irony.

How does one get access to the Internet? Well -- if you don't have a computer and a modem, get one. Your computer can act as a terminal, and you can use an ordinary telephone line to connect to an Internet-linked machine. These slower and simpler adjuncts to the Internet can provide you with the netnews discussion groups and your own e-mail address. These are services worth having -- though if you only have mail and news, you're not actually "on the Internet" proper.

If you're on a campus, your university may have direct "dedicated access" to high-speed Internet TCP/IP lines. Apply for an Internet account on a dedicated campus machine, and you may be able to get those hot-dog long-distance computing and file-transfer functions. Some cities, such as Cleveland, supply "freenet" community access. Businesses increasingly have Internet access, and are willing to sell it to subscribers. The standard fee is about $40 a month -- about the same as TV cable service.

As the Nineties proceed, finding a link to the Internet will become much cheaper and easier. Its ease of use will also improve, which is fine news, for the savage UNIX interface of TCP/IP leaves plenty of room for advancements in user-friendliness. Learning the Internet now, or at least learning about it, is wise. By the turn of the century, "network literacy," like "computer literacy" before it, will be forcing itself into the very texture of your life.

For Further Reading:

The Whole Internet Catalog & User's Guide by Ed Krol. (1992) O'Reilly and Associates, Inc.
A clear, non-jargonized introduction to the intimidating business of network literacy. Many computer-documentation manuals attempt to be funny. Mr. Krol's book is *actually* funny.

The Matrix: Computer Networks and Conferencing Systems Worldwide. by John Quarterman. Digital Press: Bedford, MA. (1990)
Massive and highly technical compendium detailing the mind-boggling scope and complexity of our newly networked planet.

The Internet Companion by Tracy LaQuey with Jeanne C. Ryer (1992) Addison Wesley.
Evangelical etiquette guide to the Internet featuring anecdotal tales of life-changing Internet experiences. Foreword by Senator Al Gore.

Zen and the Art of the Internet: A Beginner's Guide by Brendan P. Kehoe (1992) Prentice Hall.
Brief but useful Internet guide with plenty of good advice on useful machines to paw over for data. Mr Kehoe's guide bears the singularly wonderful distinction of being available in electronic form free of charge.

I'm doing the same with all my F&SF Science articles, including, of course, this one.

[end]