Chapter 2: Background to the Internet

"The revolution will not be televised, but the proceedings will be available online." Quoted from *Whole Earth Review,* Spring 1994, p.27.

The Internet is a global resource connecting millions of users; it began as an experiment over 20 years ago by the United States; specifically, its Department of Defense.' Despite its militaristic roots, today's environment of the Internet was envisioned some twenty five years ago by J.C.R. Licklider and Robert Taylor in 'The Computer as a Communication Device' *(Science and Technology,* April 1968). Licklider brought to his leadership of the Department of Defense's scientific and research network a vision of "the intergalactic computer network." Then in his essay 'The Computer as a Communication Device' Dr. Licklider foresaw the impact of broad electronic access to low cost information:

> "But to communicate is more than to send and to receive . . We believe that communicators have to do something non-trivial with the information they send and receive. And to interact with the richness of living information — not merely in the passive way that we have become accustomed to using books and libraries, but as active participants in an ongoing process, bringing something to it through our interaction with it, and not simply receiving from it by our connection to it. . . We want to emphasise something beyond its one-way transfer: the increasing significance of the jointly constructive, the mutually reinforcing aspect of communication — the part that transcends 'now we both know a fact that only one of us knew before.' *When minds interact, new ideas emerge.* " [Emphasis added.]

The phenomenal surge of interest in the Internet by for-profit firms, small businesses and individuals worldwide, is unprecedented in the electronic information industry. What the popular press and those unfamiliar with the Internet do not always appreciate is that the Internet is based on different principles from purely commercial online services. Furthermore, the vectors of change, propelled by use of the Internet, are moving rapidly and in many dimensions. The Internet is, of course, a new information environment. It is cyberspace. To understand the challenges, the opportunities, even the rhetoric of the Internet, it is important to understand where it comes from and how the climate of the Internet's information environment has evolved over time. Although the past cannot predict the future, for the Internet the past has nurtured an electronic environment that is inventing and sustaining itself. Unlike commercial online services that have a focus and a

[1] When *inter-net* is used with a small *i*, the reference may be to any network built upon TCP/IP. Small *i* internet may also refer to networks using other protocol families that can provide data in the IP format. The use of the term internet does not necessarily mean that the network is a part of the Internet or IP compliant.

finite amount of information available, the Internet is an amalgam of many electronic information flows. It spawns information by its very presence. In a sense, it is without limits. If limits exist, each day the boundaries are pushed and stretched by technology. The ultimate shape of the Internet is, therefore, difficult to define with precision. The Internet is an American invention. But its impact is global. Internet is a glimpse of the future of networked computing, a Total Network in which a user can glide seamlessly from network to network as need dictates.

1. The origin of the Internet

The originating network was named ARPANET and designed to support military research into computer communications. From its inception, ARPANET was built on a key assumption: the network is unreliable. Translating this into more practical English, the network was designed to operate during a nuclear attack as a means of allowing data to find their ways to their destinations and to rely upon many computers in many places. A nuclear blast or routine network outage would not impair the webs of communication paths.

The core of the design was a computer that would act as a switch to route the packets of data back and forth among their sources and destinations. The model for the design was somewhat like the American Telephone & Telegraph company's telephone system in the United States. Each computer was, and is still today, connected to a local switch from which all other computers can be contacted. The designers engineered additional features into the packet scheme to make certain that data arrived as they were sent.

Almost as a by-product of reliability was a foreshadowing of what is now called client-server architecture: in effect, a larger machine linked to smaller machines which could in turn be linked to other computers. Communication took place between the source and destination computer. The network required a computer to accept messages or data packets and keep the pipeline filled. If one of these traffic-cop computers failed, the rest of the network was not affected.

2. Benefits

There were, in retrospect, numerous benefits to this approach. First of all, this design required minimal information from computer clients; for a client to send a message, the computer simply put the data into an Internet Protocol (IP) packet, and 'addressed' the packets correctly. Secondly, the network architecture was a pipe through which information flowed. Thirdly, the systems environment could accommodate a wide range of computer hardware and communications devices; the Internet Protocol provided a common language, a digital Esperanto that allowed proprietary systems to make a connection if the systems could send and receive data in the IP packet format. Finally, the communicating computers — not the physical network itself — were responsible for ensuring that the communication was completed successfully. The original crucial insight was recognising that every computer on the network should be able to communicate on an equal basis with any and every other computer.

Since one of the primary goals of the network architecture was overall network survivability, the packet switches were designed to switch from one circuit to another in the event that any given circuit became congested or was interrupted. Another novelty was the introduction of a suite of protocols that could be programmed into computers connected to the ARPANET. These protocols would make it possible to transmit packets over a network composed of diverse physical media and circuits of different bit rates. By the 1980s, these protocols had evolved and allowed multiple and diverse networks to be connected to each other and thus to provide end-to-end service across many different networks. These mature protocols were called Transmission Control Protocol and Internet Protocol (TCP/IP).

3. Drawbacks

There were, of course, drawbacks associated with the approach. Firstly, because it was government funded, many large computer and telecommunications companies supported the ARPANET just because of the government funds that were poured into the programme. Secondly, expertise in using and enhancing the system did not spread much beyond major universities, research laboratories and various US government agencies. Thirdly, the commercial market pursued proprietary systems, and when the personal computer revolution began, focused more and more attention on that market. As a result, supporting, creating products for the Internet, and exploiting Internet Protocols as a solution to certain thorny networking problems was until recently a secondary or tertiary interest of many high-profile companies. In the shadows of what many companies believed was mainstream computing, a new, robust, innovative infrastructure of developers and customers was incubating. Within the last twelve months, the world has discovered what a few have known for almost 30 years: the Internet is the future of online information.

4. The emergence of the community

Perhaps the single most important realisation of the ARPANET by the mid-1970s was that a community of different computers and operating systems could communicate with each other. At first the ARPANET grew slowly, but in the 1970s it added one new computer every twenty days. By the early 1980s, the ARPANET was acquiring an increasing number of military sites, and it became clear that for security purposes there would have to be a split between research and military use. Thus MILNET (the military network) was created and diverged from ARPANET. This was a tribute to the success of the ARPANET, but it also called into question how ARPANET's future would be funded, once the Defense Department had gone its separate networking way.

5. Today's Internet-an unexpected entity

After the split between ARPANET and MILNET, the name *Internet* entered the community's vocabulary for the network referent. Grave concerns grew over the funding issue, and various schemes were advanced for managing the Internet. Fortunately for the Internet community, in the early 1980s the National Science Foundation (NSF) had elevated supercomputing to a national science priority. Five supercomputer centres were established around the United States, and NSF funded further growth and expansion of the Internet as a means of enabling users remote from any of the five supercomputing sites to have access to supercomputing.

The challenge then was to increase dramatically the speed of the network from a maximum speed of 56 kilobits per second to 1.5 megabits per second. Many in the community felt that this 28-fold increase in network speed would defeat the TCP/IP protocols, but this proved false, and now some information highway proponents are lobbying for speeds from three to five gigabits per second by the year 2000.

In retrospect, the designers of the original ARPANET made a decision that has proven itself through time. As such bodies as the Organisation for International Standardisation (ISO) struggled with a definition for a computer networking standard, the ARPANET approach became the practical way to link different machines into a working interactive network. Developers who specialised in Internet Protocol software in the United States and Western Europe offered Internet products for UNIX and virtually all other computer platforms. Over time, the collection of linked networks — likened to an electronic web — expanded rapidly and almost invisibly.

In effect, the personal computer explosion attracted the attention of computer novices and the media turned its attention upon the exciting stories of business superstars such as Microsoft and the dramatic rise and fall of the fortunes associated with the personal computer revolution.

At the same time, the equally dramatic growth of online services that became lumped under the collective name *Internet* continued. It is only within the last 12 to 18 months that the popular press, business and the general population have discovered what has been to government, academic and research professionals one of the most important computer developments ever.

In the shadows of the more newsworthy stories that filled the pages of popular computing magazines such as *PC Magazine* and the business pages of the *The Wall Street Journal,* Internet developers in the United States, United Kingdom and Scandinavia, responding to market pressures, began to put their IP software on every conceivable type of computer. It became the only practical method for computers from different manufacturers to communicate. This was attractive to the government and universities, which did not have policies saying that all computers must be bought from the same vendor. Everyone bought whichever computer they liked, and expected the computers to work together over the network.

6. A working definition

It is understandable that when the question ‘What *is* the Internet?’ is asked, the answers often vary widely. For example, representative answers may include:

- a network of networks based on the TCP/IP protocols
- a community of people who use and develop those networks
- a collection of resources that can be reached from those networks.

There are a number of reasons for the divergence of opinion about what the Internet is. Firstly, it can be thought about in several different ways; for example:

- its common protocols
- a physical collection of routers and circuits
- a set of shared resources
- an open approach to electronic interaction,

Secondly, the definitions usually reflect particular experiences with some specific aspect of a wide range of services available from a computer network. The individual's impressions are useful, but limited.

Thirdly, the Internet is a different type of computing environment. It is one that is invented as it is used. Because there is no single point of control, the Internet is an environment, not a single software package nor a specific collection of applications. The environment of the Internet does not lend itself to a simple description. When encountered by someone unfamiliar with it, the reactions range from the awe struck ('How could I have been unaware of this information resource?') to the naive ('the Internet is too disorganised and complicated to be a serious threat to commercial online services.') Both extremes are correct — up to a point.

The slippery nature of the Internet is due in part to the fact that the networks that make up the Internet are based on a standard set of protocols (a mutually agreed upon method of communication between parties) usually described as TCP/IP. (For a fuller explanation of this term, see the Glossary.) However, the Internet also has gateways to networks and services that are based on other protocols. A map of the system, if it could be drawn, would resemble a terrarium filled with a rich, complex, multi-dimensional living environment.

The popular computer press tracked the development of the IBM PC and its rapid deployment through businesses in the United States and Western Europe because it sold advertisements and was easily understood. At about the same time as the personal computer was capturing headlines, Ethernet local area networks (LANs) were beginning to find their way into government, academic and research computing environments. This technology expanded rapidly and began to mature. Powerful desktop workstations became available around 1983. Most of these workstations came with Berkeley UNIX, which included the IP networking software.

Instead of connecting one or more workstations to a single large timesharing computer per site, organisations had the tools needed to plug their local area network into the Internet. Once linked to the net, the workstations at one location could communicate with any other workstation connected anywhere in the world.

As professionals moved from one institution and organisation to another, awareness and expertise migrated as well. Simultaneously, many other organisations built their own networks using IP and its relatives. It became obvious that if these networks could talk together, users on one network could communicate with those on another; everyone would benefit. Networking facilitates the flow of information. Information by itself is useless; it only becomes useful when it is communicated and leads to the production of goods and services. Information acquires value when it gets into the right hands and is acted upon, when it moves from passive to active.

7. Mom recent initiatives

Fuelled by a growing understanding of the power of electronic information flowing among the organisations linked to the ARPANET, the United States government funded the creation of the NSFNET, commissioned by the National Science Foundation (NSF). The NSF had established five supercomputer centres. Prior to this strategic development, the world's fastest computers had only been available to certain defence contractor weapons developers and a few researchers from very large corporations. By creating supercomputer centres, the NSF was making these high-priced resources available for a broader range of defence and non-defence research. The need to provide shared access to high-performance computers placed sharply increased demands on the ARPANET for communications resources. For a variety of reasons, NSF embarked on a programme to build its own network based on the IP technology.

The NSFNET linked the five centres with 56,000 bit per second (usually abbreviated to 56 Kb/s) telecommunication lines. In the mid-1980s, the ability to transfer two typed pages of text in one second was considered a high-speed communications line. The limitations of the initial network were that only a handful of research centres and universities could be linked to the supercomputing centres because of capacity and cost constraints.

The solution adopted by the NSF was to create regional networks. In each region, institutions were connected to their nearest neighbour. Each regional cluster was connected to a supercomputer centre at one point, and the centres were connected together. This configuration allowed any computer to communicate with any other by forwarding the conversation.

This solution worked. The institutions' staffs had access to the supercomputers and to a world of data and collaborators. The network traffic increased rapidly. Merit Network, Inc., which ran Michigan's educational network in partnership with IBM and MCI, won the contract to upgrade and manage the network in 1987. The old network was replaced with faster computers and data lines that had about 20 times the capacity of the original NSFNET lines.

Over the years, the networks we group together under the rubric *Internet* have run out of capacity and been upgraded numerous times. The benefit has been the development of a network architecture that is stable, reliable and practical. More importantly, the IP network is open in the sense that the network architecture, the operating system and most of the software, are documented and publicly available.

Perhaps the most important aspect of the NSF's work is that it allowed anyone with a university or research affiliation to access the network without paying the type of fees that were associated with commercial online services that were expanding rapidly at around the same period. In the early years of the Internet, access had been available to researchers in computer science, government employees and certain government contractors.

The NSF supported educational access by funding campus connections only if the campus had a plan to provide broad access. The idea was to make it possible for anyone attending a four-year college in the United States to have access to the network. By the early 1990s, most four-year colleges are connected, and those who have experienced the power of the Internet connection are working to have secondary and primary schools connected. Over the years, people who have become Internet users have worked to have their employers obtain a connection.

There is, therefore, no specific date of birth for what is called the Internet. We can point to certain milestones; for example, the start of ARPANET and the expansion of ARPANET to MILNET in 1982. By the mid- 1980s, the Internet as most of the people in the information world now think of it, came into more visible existence.

It is important to restate that what is called the Internet is a *collection* of networks that run the TCP/IP protocols. The networks are tied together so that users of any of the networks can use the network services provided by TCP/IP to reach users on any of the other networks. This includes such networks as NSFNET, NYSERnet, and thousands of others. There are other important wide area networks, such as BITNET and DECnet which are not based on the TCP/IP protocols and are, therefore, not part of the Internet. However, it is possible to communicate between them and the Internet via electronic mail because of mail gateways that act as translators between the different network protocols involved.

Today the Internet has connections in the United States, Canada, Mexico, Central and South America, Europe, the former Soviet Union and client states, Australasia, the Pacific Basin and Asia. Currently the Internet has links in most developed countries, consists of tens of thousands of sites, and has gateways to other networks such as BITNET, JANET and EUnet that make connecting all across the world virtually seamless.

Despite the talk about commercialising the Internet, it is unlikely that a US government funded network that is freely available to qualified users will disappear. In the event that today's Internet becomes a fully commercial entity, the US defence establishment will fund and operate its own secure system. The pattern will be to build parallel, separate networks. Access to the secure network would be tightly controlled. Users of the secure network would have access to the Internet, but only authorised Internet users would have access to the secure network. When more connections are made with telephony and television data streams, the US will approach a Total Network. This will take years, but will continue until most developed countries achieve a similar connective state.

8. Who runs the Internet?

The short answer is, 'No one.' A more elaborate answer is a loose federation of US government agencies, trade associations, contractors, volunteers and committees. In its early stages, the Internet consisted of government-sponsored networks. With the advent of Xerox PARC's Ethernet, however, linking computers accelerated. Over time, privately owned and operated networks became an important part of the Internet architecture.

Organisational Structure of the Internet

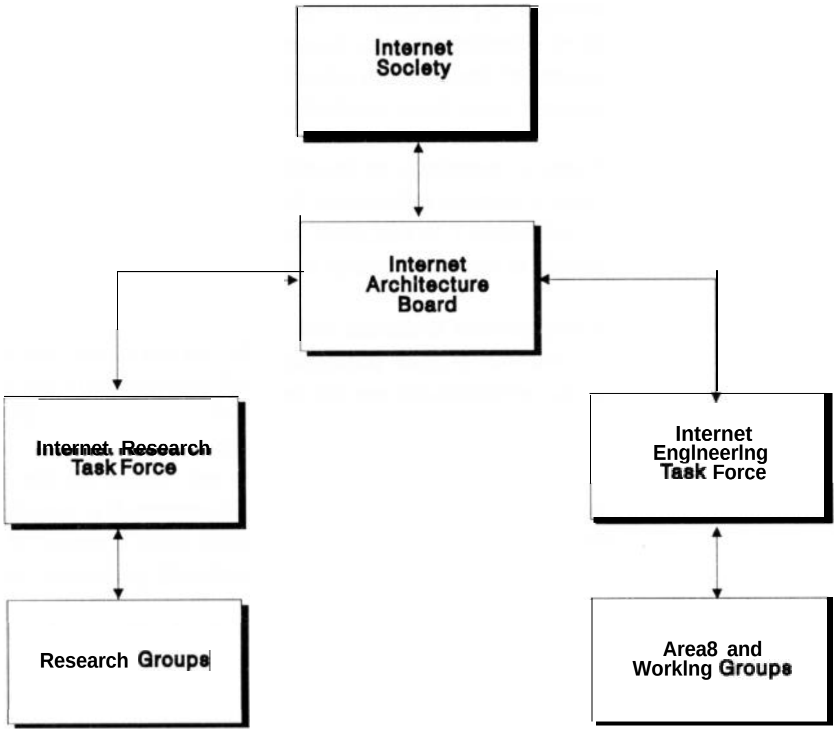

With the introduction of the National Science Foundation NSFNET, however, the architecture evolved to include intermediate-level networks consisting of collections of commercially-produced routers and trunk or access lines which connected local area network facilities to the government-sponsored backbones.

The government-sponsored supercomputer centres (such as the National Aerospace Simulator at NASA/AMES, the Magnetic Fusion Energy Computing Center at Lawrence Livermore Laboratory and the half-dozen or so NSF-sponsored supercomputer centres) fostered the growth of communications networks specifically to support supercomputer access although, over time, these have tended to look more and more like general-purpose intermediate-level networks.

A number of intermediate-level networks applied for and received funding from the National Science Foundation. These direct subsidies are temporary in the sense that those receiving funding should become self supporting. To achieve this goal, intermediate-level networks have been developing ways to generate revenue.

But the mix of government backbones, consortium intermediate level nets and private local area networks, must evolve to respond to the global, economic and technical demands of the user communities. With the move to higher speed backbone operation, the importance of the participation in the Internet of a broader range of non-governmental organisations has increased. However, the US government funding for Internet is going to decrease. The Internet must become self-supporting. In other words, it is in the process of changing from a government-sponsored service with a user community composed of academics, researchers and a handful of commercial activities, into a different, marketplace-driven environment.

9. As the twig is bent

As we have seen, the entity now labelled 'the Internet' once consisted solely of government-sponsored networks such as the Defense Department's ARPANET and a handful of wireless links; namely, packet radio networks and packet satellite networks.

Network technology provided the catalyst for rapid evolution which continues to this time. The introduction of Xerox PARC's Ethernet more than 15 years ago produced the second wave of innovation. Today's high-speed networks are building momentum for a third wave of development and its corresponding explosion of innovation.

An increasing number of privately owned and operated networks will remain an essential feature of the Internet architecture. As commercialisation gains momentum, a mixture of government-sponsored backbone facilities and private local area networks will comprise the Internet networking environment. When the National Science Foundation introduced its NSFNET, intermediate-level networks sprang into existence. These consisted of collections of commercially-produced routers and trunk or access lines which connected local area network facilities to the government-sponsored backbones.

10. American as apple pie

The Internet is a peculiarly American institution. It blended from the outset what many countries would consider proprietary resources with loosely-regulated access. Thus, government-sponsored supercomputer centres fostered the growth of communications networks specifically to support supercomputer access. In addition, the network created an electronic community of those who had public access to the Internet and discovered others who shared access and interests. Today, the blend of networks has become a general purpose network offering a broad range of information and services. In short, Internet has become a public supermarket while remaining a high technology boutique.

What are known as intermediate-level networks like those at most major universities received seed funding from the National Science Foundation. The expectation exists that such subsidies will continue despite the increasingly vociferous support for cost recovery and self funding. Decades ago, the National Science Foundation took the position that these intermediate networks should become self-sustaining. To accomplish this objective, at least in part, most intermediate-level networks have been stepping up their marketing efforts.

The basic model is a hotchpotch of US government-supported high-speed backbones, federations of intermediate level nets, and private local area networks,

works. However, the surge in usage and the demand for larger and larger communications bandwidth sets the stage for change. There is, therefore, little chance that the commercialisation trend will diminish.

The high-speed gigabit systems will be built and operated by commercial entities. Users who want access to the more advanced network will have to pay for it. If there is Federal subsidy in this model, US government support is likely to come through direct provision of support for networking at the level of individual research grant or possibly through a system of institutional vouchers permitting and perhaps even mandating institution-wide network planning and provision. This differs from the present model in which the backbone networks are essentially federally owned and operated, or enjoy significant, direct federal support to the provider of the service.

11. Outlook 2000

The importance of such a shift in service provision philosophy cannot be ignored. It promises to eliminate restrictions on the use and application of the backbone facilities, changes free access into access by those who can pay fees set by commercial providers.

In the international arena, private networks ranging from local consortia to intercontinental systems will be operated by a range of entities. The economics of private networking may still be favourable for sufficiently heavy usage.

The stage is therefore set for a chess game with many players acting out their strategies on a global scale. Competitive pressures and political considerations will play an increasingly large part in the Internet's development. Despite the best efforts of nation states, the Internet will continue on its revolutionary path.