

What Is The Internet (And What Makes It Work) - December, 1999  
By Robert E. Kahn and Vinton G. Cerf

In this issue:

INTRODUCTION  
THE EVOLUTION OF THE INTERNET  
THE INTERNET ARCHITECTURE  
GOVERNMENT'S HISTORICAL ROLE  
A DEFINITION FOR THE INTERNET  
WHO RUNS THE INTERNET  
WHERE DO WE GO FROM HERE?

This paper was prepared by the authors at the request of the Internet Policy Institute (IPI), a non-profit organization based in Washington, D.C., for inclusion in their upcoming series of Internet related papers. It is a condensation of a longer paper in preparation by the authors on the same subject. Many topics of potential interest were not included in this condensed version because of size and subject matter constraints. Nevertheless, the reader should get a basic idea of the Internet, how it came to be, and perhaps even how to begin thinking about it from an architectural perspective. This will be especially important to policy makers who need to distinguish the Internet as a global information system apart from its underlying communications infrastructure.

## INTRODUCTION

As we approach a new millennium, the Internet is revolutionizing our society, our economy and our technological systems. No one knows for certain how far, or in what direction, the Internet will evolve. But no one should underestimate its importance.

Over the past century and a half, important technological developments have created a global environment that is drawing the people of the world closer and closer together. During the industrial revolution, we learned to put motors to work to magnify human and animal muscle power. In the new Information Age, we are learning to magnify brainpower by putting the power

of computation wherever we need it, and to provide information services on a global basis. Computer resources are infinitely flexible tools; networked together, they allow us to generate, exchange, share and manipulate information in an uncountable number of ways. The Internet, as an integrating force, has melded the technology of communications and computing to provide instant connectivity and global information services to all its users at very low cost.

Ten years ago, most of the world knew little or nothing about the Internet. It was the private enclave of computer scientists and researchers who used it to interact with colleagues in their respective disciplines. Today, the Internet's magnitude is thousands of times what it was only a decade ago. It is estimated that about 60 million host computers on the Internet today serve about 200 million users in over 200 countries and territories. Today's telephone system is still much larger: about 3 billion people around the world now talk on almost 950 million telephone lines (about 250 million of which are actually radio-based cell phones). But by the end of the year 2000, the authors estimate there will be at least 300 million Internet users. Also, the total numbers of host computers and users have been growing at about 33% every six months since 1988 — or roughly 80% per year. The telephone service, in comparison, grows an average of about 5-10% per year. That means if the Internet keeps growing steadily the way it has been growing over the past few years, it will be nearly as big as today's telephone system by about 2006.

## THE EVOLUTION OF THE INTERNET

The underpinnings of the Internet are formed by the global interconnection of hundreds of thousands of otherwise independent computers, communications entities and information systems. What makes this interconnection possible is the use of a set of communication standards, procedures and formats in common among the networks and the various devices and computational facilities connected to them. The procedures by which computers communicate with each other are called "protocols." While this infrastructure is steadily evolving to include new capabilities, the protocols initially used by the Internet are called the "TCP/IP" protocols, named after the two protocols that formed the principal basis for Internet operation.

On top of this infrastructure is an emerging set of architectural concepts and data structures for heterogeneous information systems that renders the Internet a truly global information system. In essence, the Internet is an architecture, although many people confuse it with its implementation. When the Internet is looked at as an architecture, it manifests two different abstractions. One abstraction deals with communications connectivity, packet delivery and a variety of end-end communication services. The other abstraction deals with the Internet as an information system, independent of its underlying communications infrastructure, which allows creation, storage and access to a wide range of information resources, including digital objects and related services at various levels of abstraction.

Interconnecting computers is an inherently digital problem. Computers process and exchange digital information, meaning that they use a discrete mathematical binary or two-valued language of 1s and 0s. For communication purposes, such information is mapped into continuous electrical or optical waveforms. The use of digital signaling allows accurate regeneration and reliable recovery of the underlying bits. We use the terms computer, computer resources and computation to mean not only traditional computers, but also devices that can be controlled digitally over a network, information resources such as mobile programs and other computational capabilities.

The telephone network started out with operators who manually connected telephones to each other through patch panels that accepted patch cords from each telephone line and electrically connected them to one another through the panel, which operated, in effect, like a switch. The result was called circuit switching, since at its conclusion, an electrical circuit was made between the calling telephone and the called telephone. Conventional circuit switching, which was developed to handle telephone calls, is inappropriate for connecting computers because it makes limited use of the telecommunication facilities and takes too long to set up connections. Although reliable enough for voice communication, the circuit-switched voice network had difficulty delivering digital information without errors.

For digital communications, packet switching is a better choice, because it is far better suited to the typically "burst" communication style of computers. Computers that communicate typically send out brief but intense bursts of data, then remain silent for a while before sending out the

next burst. These bursts are communicated as packets, which are very much like electronic postcards. The postcards, in reality packets, are relayed from computer to computer until they reach their destination. The special computers that perform this forwarding function are called variously "packet switches" or "routers" and form the equivalent of many bucket brigades spanning continents and oceans, moving buckets of electronic postcards from one computer to another. Together these routers and the communication links between them form the underpinnings of the Internet.

Without packet switching, the Internet would not exist as we now know it. Going back to the postcard analogy, postcards can get lost. They can be delivered out of order, and they can be delayed by varying amounts. The same is true of Internet packets, which, on the Internet, can even be duplicated. The Internet Protocol is the postcard layer of the Internet. The next higher layer of protocol, TCP, takes care of re-sending the postcards to recover packets that might have been lost, and putting packets back in order if they have become disordered in transit.

Of course, packet switching is about a billion times faster than the postal service or a bucket brigade would be. It also has to operate over many different communications systems, or substrata. The authors designed the basic architecture to be so simple and undemanding that it could work with most communication services. Many organizations, including commercial ones, carried out research using the TCP/IP protocols in the 1970s. Email was steadily used over the nascent Internet during that time and to the present. It was not until 1994 that the general public began to be aware of the Internet by way of the World Wide Web application, particularly after Netscape Communications was formed and released its browser and associated server software.

Thus, the evolution of the Internet was based on two technologies and a research dream. The technologies were packet switching and computer technology, which, in turn, drew upon the underlying technologies of digital communications and semiconductors. The research dream was to share information and computational resources. But that is simply the technical side of the story. Equally important in many ways were the other dimensions that enabled the Internet to come into existence and flourish. This aspect of the story starts with cooperation and far-sightedness in the U.S. Government, which is often derided for lack of foresight but is a real hero in this story.

It leads on to the enthusiasm of private sector interests to build upon the government funded developments to expand the Internet and make it available to the general public. Perhaps most important, it is fueled by the development of the personal computer industry and significant changes in the telecommunications industry in the 1980s, not the least of which was the decision to open the long distance market to competition. The role of workstations, the Unix operating system and local area networking (especially the Ethernet) are themes contributing to the spread of Internet technology in the 1980s into the research and academic community from which the Internet industry eventually emerged.

Many individuals have been involved in the development and evolution of the Internet covering a span of almost four decades if one goes back to the early writings on the subject of computer networking by Kleinrock [i], Licklider [ii], Baran [iii], Roberts [iv], and Davies [v]. The ARPANET, described below, was the first wide-area computer network. The NSFNET, which followed more than a decade later under the leadership of Erich Bloch, Gordon Bell, Bill Wulf and Steve Wolff, brought computer networking into the mainstream of the research and education communities. It is not our intent here to attempt to attribute credit to all those whose contributions were central to this story, although we mention a few of the key players. A readable summary on the history of the Internet, written by many of the key players, may be found at [www.isoc.org/internet/history](http://www.isoc.org/internet/history). [vi]

### From One Network to Many: The role of DARPA

Modern computer networking technologies emerged in the early 1970s. In 1969, The U.S. Defense Advanced Research Projects Agency (variously called ARPA and DARPA), an agency within the Department of Defense, commissioned a wide-area computer network called the ARPANET. This network made use of the new packet switching concepts for interconnecting computers and initially linked computers at universities and other research institutions in the United States and in selected NATO countries. At that time, the ARPANET was essentially the only realistic wide-area computer network in existence, with a base of several dozen organizations, perhaps twice that number of computers and numerous researchers at those sites. The program was led at DARPA by Larry Roberts. The packet switches

were built by Bolt Beranek and Newman (BBN), a DARPA contractor. Others directly involved in the ARPANET activity included the authors, Len Kleinrock, Frank Heart, Howard Frank, Steve Crocker, Jon Postel and many many others in the ARPA research community.

Back then, the methods of internetworking (that is interconnecting computer networks) were primitive or non-existent. Two organizations could interwork technically by agreeing to use common equipment, but not every organization was interested in this approach. Absent that, there was jury-rigging, special case development and not much else. Each of these networks stood on its own with essentially no interaction between them a far cry from today's Internet.

In the early 1970s, ARPA began to explore two alternative applications of packet switching technology based on the use of synchronous satellites (SATNET) and ground-based packet radio (PRNET). The decision by Kahn to link these two networks and the ARPANET as separate and independent networks resulted in the creation of the Internet program and the subsequent collaboration with Cerf. These two systems differed in significant ways from the ARPANET so as to take advantage of the broadcast and wireless aspects of radio communications. The strategy that had been adopted for SATNET originally was to embed the SATNET software into an ARPANET packet switch, and interwork the two networks through memory-to-memory transfers within the packet switch. This approach, in place at the time, was to make SATNET an embedded network within the ARPANET; users of the network would not even need to know of its existence. The technical team at Bolt Beranek and Newman (BBN), having built the ARPANET switches and now building the SATNET software, could easily produce the necessary patches to glue the programs together in the same machine. Indeed, this is what they were under contract with DARPA to provide. By embedding each new network into the ARPANET, a seamless internetworked capability was possible, but with no realistic possibility of unleashing the entrepreneurial networking spirit that has manifest itself in modern day Internet developments. A new approach was in order.

The Packet Radio (PRNET) program had not yet gotten underway so there was ample opportunity to change the approach there. In addition, up until then, the SATNET program was only an equipment development activity. No commitments had been obtained for the use of actual satellites or ground stations to access them. Indeed, since there was no domestic satellite industry in the U.S. then, the only two viable alternatives were the use of Intelsat or U.S. military satellites. The time for a change in strategy, if it was to be made, was then.

## THE INTERNET ARCHITECTURE

The authors created an architecture for interconnecting independent networks that could then be federated into a seamless whole without changing any of the underlying networks. This was the genesis of the Internet as we know it today.

In order to work properly, the architecture required a global addressing mechanism (or Internet address) to enable computers on any network to reference and communicate with computers on any other network in the federation. Internet addresses fill essentially the same role as telephone numbers do in telephone networks. The design of the Internet assumed first that the individual networks could not be changed to accommodate new architectural requirements; but this was largely a pragmatic assumption to facilitate progress. The networks also had varying degrees of reliability and speed. Host computers would have to be able to put disordered packets back into the correct order and discard duplicate packets that had been generated along the way. This was a major change from the virtual circuit-like service provided by ARPANET and by then contemporary commercial data networking services such as Tymnet and Telenet. In these networks, the underlying network took responsibility for keeping all information in order and for re-sending any data that might have been lost. The Internet design made the computers responsible for tending to these network problems.

A key architectural construct was the introduction of gateways (now called routers) between the networks to handle the disparities such as different data rates, packet sizes, error conditions, and interface specifications. The gateways would also check the destination Internet addresses of each

packet to determine the gateway to which it should be forwarded. These functions would be combined with certain end-end functions to produce the reliable communication from source to destination. A draft paper by the authors describing this approach was given at a meeting of the International Network Working Group in 1973 in Sussex, England and the final paper was subsequently published by the Institute for Electrical and Electronics Engineers, the leading professional society for the electrical engineering profession, in its Transactions on Communications in May, 1974 [vii]. The paper described the TCP/IP protocol.

DARPA contracted with Cerf's group at Stanford to carry out the initial detailed design of the TCP software and, shortly thereafter, with BBN and University College London to build independent implementations of the TCP protocol (as it was then called – it was later split into TCP and IP) for different machines. BBN also had a contract to build a prototype version of the gateway. These three sites collaborated in the development and testing of the initial protocols on different machines. Cerf, then a professor at Stanford, provided the day-to-day leadership in the initial TCP software design and testing. BBN deployed the gateways between the ARPANET and the PRNET and also with SATNET. During this period, under Kahn's overall leadership at DARPA, the initial feasibility of the Internet Architecture was demonstrated.

The TCP/IP protocol suite was developed and refined over a period of four more years and, in 1980, it was adopted as a standard by the U.S. Department of Defense. On January 1, 1983 the ARPANET converted to TCP/IP as its standard host protocol. Gateways (or routers) were used to pass packets to and from host computers on local area networks. Refinement and extension of these protocols and many others associated with them continues to this day by way of the Internet Engineering Task Force [viii].

## GOVERNMENT'S HISTORICAL ROLE

Other political and social dimensions that enabled the Internet to come into existence and flourish are just as important as the technology upon which it is based. The federal government played a large role in creating the Internet, as did the private sector interests that made it available to the general



public. The development of the personal computer industry and significant changes in the telecommunications industry also contributed to the Internet's growth in the 1980s. In particular, the development of workstations, the Unix operating system, and local area networking (especially the Ethernet) contributed to the spread of the Internet within the research community from which the Internet industry eventually emerged.

The National Science Foundation and others

In the late 1970s, the National Science Foundation (NSF) became interested in the impact of the ARPANET on computer science and engineering. NSF funded the Computer Science Network (CSNET), which was a logical design for interconnecting universities that were already on the ARPANET and those that were not. Telenet was used for sites not connected directly to the ARPANET and a gateway was provided to link the two. Independent of NSF, another initiative called BITNET ("Because it's there" Net) [ix] provided campus computers with email connections to the growing ARPANET. Finally, AT&T Bell Laboratories development of the Unix operating system led to the creation of a grass-roots network called USENET [x], which rapidly became home to thousands of newsgroups where Internet users discussed everything from aerobics to politics and zoology.

In the mid 1980s, NSF decided to build a network called NSFNET to provide better computer connections for the science and education communities. The NSFNET made possible the involvement of a large segment of the education and research community in the use of high speed networks. A consortium consisting of MERIT (a University of Michigan non-profit network services organization), IBM and MCI Communications won a 1987 competition for the contract to handle the network's construction. Within two years, the newly expanded NSFNET had become the primary backbone component of the Internet, augmenting the ARPANET until it was decommissioned in 1990. At about the same time, other parts of the U.S. government had moved ahead to build and deploy networks of their own, including NASA and the Department of Energy. While these groups originally adopted independent approaches for their networks, they

eventually decided to support the use of TCP/IP.

The developers of the NSFNET, led by Steve Wolff who had the direct responsibility for the NSFNET program, also decided to create intermediate level networks to serve research and education institutions and, more importantly, to allow networks that were not commissioned by the U.S. government to connect to the NSFNET. Nearly a dozen intermediate level networks were created, most with NSF support. This strategy reduced the overall load on the backbone network operators and spawned a new industry: Internet Service Provision. Nearly a dozen intermediate level networks were created, most with NSF support, [xi] some, such as UUNET, with Defense support, and some without any government support. The NSF contribution to the evolution of the Internet was essential in two respects. It opened the Internet to many new users and, drawing on the properties of TCP/IP, structured it so as to allow many more network service providers to participate.

For a long time, the federal government did not allow organizations to connect to the Internet to carry out commercial activities. By 1988, it was becoming apparent, however, that the Internet's growth and use in the business sector might be seriously inhibited by this restriction. That year, CNRI requested permission from the Federal Networking Council to interconnect the commercial MCI Mail electronic mail system to the Internet as part of a general electronic mail interconnection experiment. Permission was given and the interconnection was completed by CNRI, under Cerf's direction, in the summer of 1989. Shortly thereafter, two of the then non-profit Internet Service Providers (UUNET [xii] and NYSERNET) produced new for-profit companies (UUNET and PSINET [xiii] respectively). In 1991, they were interconnected with each other and CERFNET [xiv]. Commercial pressure to alleviate restrictions on interconnections with the NSFNET began to mount.

In response, Congress passed legislation allowing NSF to open the NSFNET to commercial usage. Shortly thereafter, NSF determined that its support for NSFNET might not be required in the longer term and, in April 1995, NSF ceased its support for the NSFNET. By that time, many commercial networks were in operation and provided alternatives to NSFNET for national level network services. Today, approximately 10,000 Internet

Service Providers (ISPs) are in operation. Roughly half the world's ISPs currently are based in North America and the rest are distributed throughout the world.

## A DEFINITION FOR THE INTERNET

The authors feel strongly that efforts should be made at top policy levels to define the Internet. It is tempting to view it merely as a collection of networks and computers. However, as indicated earlier, the authors designed the Internet as an architecture that provided for both communications capabilities and information services. Governments are passing legislation pertaining to the Internet without ever specifying to what the law applies and to what it does not apply. In U.S. telecommunications law, distinctions are made between cable, satellite broadcast and common carrier services. These and many other distinctions all blur in the backdrop of the Internet. Should broadcast stations be viewed as Internet Service Providers when their programming is made available in the Internet environment? Is use of cellular telephones considered part of the Internet and if so under what conditions? This area is badly in need of clarification.

The authors believe the best definition currently in existence is that approved by the Federal Networking Council in 1995, <http://www.fnc.gov> and which is

reproduced in the footnote below [xv] for ready reference. Of particular note is that it defines the Internet as a global information system, and included in the definition, is not only the underlying communications technology, but also higher-level protocols and end-user applications, the associated data structures and the means by which the information may be processed, manifested, or otherwise used. In many ways, this definition supports the characterization of the Internet as an information superhighway. Like the federal highway system, whose underpinnings include not only concrete lanes and on/off ramps, but also a supporting infrastructure both physical and informational, including signs, maps, regulations, and such related services and products as filling stations and gasoline, the Internet has its own layers of ingress and egress, and its own multi-tiered levels of service.

The FNC definition makes it clear that the Internet is a dynamic organism that can be looked at in myriad ways. It is a framework for numerous

services and a medium for creativity and innovation. Most importantly, it can be expected to evolve.

## WHO RUNS THE INTERNET

### The Domain Name System

The Internet evolved as an experimental system during the 1970s and early 1980s. It then flourished after the TCP/IP protocols were made mandatory on the ARPANET and other networks in January 1983; these protocols thus became the standard for many other networks as well. Indeed, the Internet grew so rapidly that the existing mechanisms for associating the names of host computers (e.g. UCLA, USC-ISI) to Internet addresses (known as IP addresses) were about to be stretched beyond acceptable engineering limits. Most of the applications in the Internet referred to the target computers by name. These names had to be translated into Internet addresses before the lower level protocols could be activated to support the application. For a time, a group at SRI International in Menlo Park, CA, called the Network Information Center (NIC), maintained a simple, machine-readable list of names and associated Internet addresses which was made available on the net. Hosts on the Internet would simply copy this list, usually daily, so as to maintain a local copy of the table. This list was called the "host.txt" file (since it was simply a text file). The list served the function in the Internet that directory services (e.g. 411 or 703-555-1212) do in the US telephone system - the translation of a name into an address.

As the Internet grew, it became harder and harder for the NIC to keep the list current. Anticipating that this problem would only get worse as the network expanded, researchers at USC Information Sciences Institute launched an effort to design a more distributed way of providing this same information. The end result was the Domain Name System (DNS) [xvi] which allowed hundreds of thousands of "name servers" to maintain small portions of a global database of information associating IP addresses with the names of computers on the Internet.

The naming structure was hierarchical in character. For example, all host computers associated with educational institutions would have names like "stanford.edu" or "ucla.edu". Specific hosts would have names like

"cs.ucla.edu" to refer to a computer in the computer science department of UCLA, for example. A special set of computers called "root servers" maintained information about the names and addresses of other servers that contained more detailed name/address associations. The designers of the DNS also developed seven generic "top level" domains, as follows:

Education - EDU  
Government - GOV  
Military - MIL  
International - INT  
Network - NET  
(non-profit) Organization - ORG  
Commercial - COM

Under this system, for example, the host name "UCLA" became "UCLA.EDU" because it was operated by an educational institution, while the host computer for "BBN" became "BBN.COM" because it was a commercial organization. Top-level domain names also were created for every country: United Kingdom names would end in .UK, while the ending .FR was created for the names of France.

The Domain Name System (DNS) was and continues to be a major element of the Internet architecture, which contributes to its scalability. It also contributes to controversy over trademarks and general rules for the creation and use of domain names, creation of new top-level domains and the like. At the same time, other resolution schemes exist as well. One of the authors (Kahn) has been involved in the development of a different kind of standard identification and resolution scheme [xvii] that, for example, is being used as the base technology by book publishers to identify books on the Internet by adapting various identification schemes for use in the Internet environment. For example, International Standard Book Numbers (ISBNs) can be used as part of the identifiers. The identifiers then resolve to state information about the referenced books, such as location information (e.g. multiple sites) on the Internet that is used to access the books or to order them. These developments are taking place in parallel with the more traditional means of managing Internet resources. They offer an alternative to the existing Domain Name System with enhanced functionality.

The growth of Web servers and users of the Web has been remarkable, but some people are confused about the relationship between the World Wide

Web and the Internet. The Internet is the global information system that includes communication capabilities and many high level applications. The Web is one such application. The existing connectivity of the Internet made it possible for users and servers all over the world to participate in this activity. Electronic mail is another important application. As of today, over 60 million computers take part in the Internet and about 3.6 million web sites were estimated to be accessible on the net. Virtually every user of the net has access to electronic mail and web browsing capability. Email remains a critically important application for most users of the Internet, and these two functions largely dominate the use of the Internet for most users.

### The Internet Standards Process

Internet standards were once the output of research activity sponsored by DARPA. The principal investigators on the internetting research effort essentially determined what technical features of the TCP/IP protocols would become common. The initial work in this area started with the joint effort of the two authors, continued in Cerf's group at Stanford, and soon thereafter was joined by engineers and scientists at BBN and University College London. This informal arrangement has changed with time and details can be found elsewhere [xviii]. At present, the standards efforts for Internet is carried out primarily under the auspices of the Internet Society (ISOC). The Internet Engineering Task Force (IETF) operates under the leadership of its Internet Engineering Steering Group (IESG), which is populated by appointees approved by the Internet Architecture Board (IAB) which is, itself, now part of the Internet Society.

The IETF comprises over one hundred working groups categorized and managed by Area Directors specializing in specific categories.

There are other bodies with considerable interest in Internet standards or in standards that must interwork with the Internet. Examples include the International Telecommunications Union Telecommunications standards group (ITU-T), the International Institute of Electrical and Electronic Engineers (IEEE) local area network standards group (IEEE 801), the Organization for International Standardization (ISO), the American National Standards Institute (ANSI), the World Wide Web Consortium (W3C), and many others.

As Internet access and services are provided by existing media such as telephone, cable and broadcast, interactions with standards bodies and legal structures formed to deal with these media will become an increasingly complex matter. The intertwining of interests is simultaneously fascinating and complicated, and has increased the need for thoughtful cooperation among many interested parties.

### Managing the Internet

Perhaps the least understood aspect of the Internet is its management. In recent years, this subject has become the subject of intense commercial and international interest, involving multiple governments and commercial organizations, and recently congressional hearings. At issue is how the Internet will be managed in the future, and, in the process, what oversight mechanisms will insure that the public interest is adequately served.

In the 1970s, managing the Internet was easy. Since few people knew about the Internet, decisions about almost everything of real policy concern were made in the offices of DARPA. It became clear in the late 1970s, however, that more community involvement in the decision-making processes was essential. In 1979, DARPA formed the Internet Configuration Control Board (ICCB) to insure that knowledgeable members of the technical community discussed critical issues, educated people outside of DARPA about the issues, and helped others to implement the TCP/IP protocols and gateway functions. At the time, there were no companies that offered turnkey solutions to getting on the Internet. It would be another five years or so before companies like Cisco Systems were formed, and while there were no PCs yet, the only workstations available were specially built and their software was not generally configured for use with external networks; they were certainly considered expensive at the time.

In 1983, the small group of roughly twelve ICCB members was reconstituted (with some substitutions) as the Internet Activities Board (IAB), and about ten Task Forces were established under it to address issues in specific technical areas. The attendees at Internet Working Group meetings were invited to become members of as many of the task forces as they wished.

The management of the Domain Name System offers a kind of microcosm of issues now frequently associated with overall management of the

Internet's operation and evolution. Someone had to take responsibility for overseeing the system's general operation. In particular, top-level domain names had to be selected, along with persons or organizations to manage each of them. Rules for the allocation of Internet addresses had to be established. DARPA had previously asked the late Jon Postel of the USC Information Sciences Institute to take on numerous functions related to administration of names, addresses and protocol related matters. With time, Postel assumed further responsibilities in this general area on his own, and DARPA, which was supporting the effort, gave its tacit approval. This activity

was generally referred to as the Internet Assigned Numbers Authority (IANA)

[xix]. In time, Postel became the arbitrator of all controversial matters concerning names and addresses until his untimely death in October 1998.

It is helpful to consider separately the problem of managing the domain name space and the Internet address space. These two vital elements of the Internet architecture have rather different characteristics that color the management problems they generate. Domain names have semantics that numbers may not imply; and thus a means of determining who can use what names is needed. As a result, speculators on Internet names often claim large numbers of them without intent to use them other than to resell them later. Alternate resolution mechanisms [xx], if widely adopted, could significantly change the landscape here.

The rapid growth of the Internet has triggered the design of a new and larger address space (the so-called IP version 6 address space); today's Internet uses IP version 4 [xxi]. However, little momentum has yet developed to deploy IPv6 widely. Despite concerns to the contrary, the IPv4 address space will not be depleted for some time. Further, the use of Dynamic Host Configuration Protocol (DHCP) to dynamically assign IP addresses has also cut down on demand for dedicated IP addresses. Nevertheless, there is growing recognition in the Internet technical community that expansion of the address space is needed, as is the development of transition schemes that allow interoperation between IPv4 and IPv6 while migrating to IPv6.

In 1998, the Internet Corporation for Assigned Names and Numbers (ICANN)

was formed as a private sector, non-profit, organization to oversee the orderly progression in use of Internet names and numbers, as well as



certain protocol related matters that required oversight. The birth of this organization, which was selected by the Department of Commerce for this function, has been difficult, embodying as it does many of the inherent conflicts in resolving discrepancies in this arena. However, there is a clear need for an oversight mechanism for Internet domain names and numbers, separate from their day-to-day management.

Many questions about Internet management remain. They may also prove difficult to resolve quickly. Of specific concern is what role the U.S. government and indeed governments around the world need to play in its continuing operation and evolution. This is clearly a subject for another time.

## WHERE DO WE GO FROM HERE?

As we struggle to envision what may be commonplace on the Internet in a decade, we are confronted with the challenge of imagining new ways of doing old things, as well as trying to think of new things that will be enabled by the Internet, and by the technologies of the future.

In the next ten years, the Internet is expected to be enormously bigger than it is today. It will be more pervasive than the older technologies and penetrate more homes than television and radio programming. Computer chips are now being built that implement the TCP/IP protocols and recently a university announced a two-chip web server. Chips like this are extremely small and cost very little. And they can be put into anything. Many of the devices connected to the Internet will be Internet-enabled appliances (cell phones, fax machines, household appliances, hand-held organizers, digital cameras, etc.) as well as traditional laptop and desktop computers. Information access will be directed to digital objects of all kinds and services that help to create them or make use of them [xxii].

Very high-speed networking has also been developing at a steady pace. From the original 50,000 bit-per-second ARPANET, to the 155 million bit-per-second NSFNET, to today's 2.4 - 9.6 billion bit-per-second commercial networks, we routinely see commercial offerings providing Internet access at increasing speeds. Experimentation with optical technology using wavelength division multiplexing is underway in many quarters; and testbeds operating at speeds of terabits per second (that is trillions of bits-per-second) are being constructed.

Some of these ultra-high speed systems may one-day carry data from very far away places, like Mars. Already, design of the interplanetary Internet as a logical extension of the current Internet, is part of the NASA Mars mission program now underway at the Jet Propulsion Laboratory in Pasadena, California [xxiii]. By 2008, we should have a well functioning Earth-Mars network that serves as a nascent backbone of the interplanetary Internet.

Wireless communication has exploded in recent years with the rapid growth of cellular telephony. Increasingly, however, Internet access is becoming available over these networks. Alternate forms for wireless communication, including both ground radio and satellite are in development and use now, and the prospects for increasing data rates look promising. Recent developments in high data rate systems appear likely to offer ubiquitous wireless data services in the 1-2 Mbps range. It is even possible that wireless Internet access may one day be the primary way most people get access to the Internet.

A developing trend that seems likely to continue in the future is an information centric view of the Internet that can live in parallel with the current communications centric view. Many of the concerns about intellectual property protection are difficult to deal with, not because of fundamental limits in the law, but rather by technological and perhaps management limitations in knowing how best to deal with these issues. A digital object infrastructure that makes information objects first-class citizens in the packetized primordial soup of the Internet is one step in that direction. In this scheme, the digital object is the conceptual elemental unit in the information view; it is interpretable (in principle) by all participating information systems. The digital object is thus an abstraction that may be implemented in various ways by different systems. It is a critical building block for interoperable and heterogeneous information systems. Each digital object has a unique and, if desired, persistent identifier that will allow it to be managed over time. This approach is highly relevant to the development of third-party value added information services in the Internet environment.

Of special concern to the authors is the need to understand and manage the downside potential for network disruptions, as well as cybercrime and terrorism. The ability to deal with problems in this diverse arena is at the forefront of maintaining a viable global information infrastructure.  
IOPS.org

[xxiv] a private-sector group dedicated to improving coordination among ISPs - deals with issues of ISP outages, disruptions, other trouble conditions, as well as related matters, by discussion, interaction and coordination between and among the principal players. Business, the academic community and government all need as much assurance as possible that they can conduct their activities on the Internet with high confidence that security and reliability will be present. The participation of many organizations around the world, including especially governments and the relevant service providers will be essential here.

The success of the Internet in society as a whole will depend less on technology than on the larger economic and social concerns that are at the heart of every major advance. The Internet is no exception, except that its potential and reach are perhaps as broad as any that have come before.

#### Footnotes

[i] Leonard Kleinrock's dissertation thesis at MIT was written during 1961: "Information Flow in Large Communication Nets", RLE Quarterly Progress Report, July 1961 and published as a book "Communication Nets: Stochastic Message Flow and Delay", New York: McGraw Hill, 1964. This was one of the earliest mathematical analyses of what we now call packet switching networks.

[ii] J.C.R. Licklider & W. Clark, "On-Line Man Computer Communication", August 1962. Licklider made tongue-in-cheek references to an "inter-galactic network" but in truth, his vision of what might be possible was prophetic.

[iii] [BARAN 64] Baran, P., et al, "On Distributed Communications", Volumes I-XI, RAND Corporation Research Documents, August 1964. Paul Baran explored the use of digital "message block" switching to support highly resilient, survivable voice communications for military command and control. This work was undertaken at RAND Corporation for the US Air Force beginning in 1962.

[iv] L. Roberts & T. Merrill, "Toward a Cooperative Network of

Time-Shared Computers", Fall AFIPS Conf., Oct. 1966.

[v] Davies, D.W., K.A. Bartlett, R.A. Scantlebury, and P. T. Wilkinson. 1967. "A Digital Communication Network for Computers Giving Rapid Response at Remote Terminals," Proceedings of the ACM Symposium on Operating System Principles. Association for Computing Machinery, New York, 1967. Donald W. Davies and his colleagues coined the term "packet" and built one node of a packet switching network at the National Physical Laboratory in the UK.

[vi] Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, Stephen Wolff, "A Brief History of the Internet," [www.isoc.org/internet/history/brief.html](http://www.isoc.org/internet/history/brief.html) and see below for timeline

[vii] Vinton G. Cerf and Robert E. Kahn, "A Protocol for Packet Network Intercommunication," IEEE Transactions on Communications, Vol. COM-22, May 1974.

[viii] The Internet Engineering Task Force (IETF) is an activity taking place under the auspices of the Internet Society ([www.isoc.org](http://www.isoc.org)). See [www.ietf.org](http://www.ietf.org)

[ix] From the BITNET charter:

BITNET, which originated in 1981 with a link between CUNY and Yale, grew rapidly during the next few years, with management and systems services provided on a volunteer basis largely from CUNY and Yale. In 1984, the BITNET Directors established an Executive Committee to provide policy guidance.

(see <http://www.geocities.com/SiliconValley/2260/bitchart.html>)

[x] Usenet came into being in late 1979, shortly after the release of V7 Unix with UUCP. Two Duke University grad students in North Carolina, Tom Truscott and Jim Ellis, thought of hooking computers together to exchange information with the Unix community. Steve Bellovin, a grad student at the University of North Carolina, put together the first version of the news software using shell scripts and installed it on the first two sites: "unc" and "duke." At the beginning of 1980 the network consisted of those two sites

and "phs" (another machine at Duke), and was described at the January Usenix conference. Steve Bellovin later rewrote the scripts into C programs, but they were never released beyond "unc" and "duke." Shortly thereafter, Steve Daniel did another implementation in C for public distribution. Tom Truscott made further modifications, and this became the "A" news release.

(see <http://www.ou.edu/research/electron/intemet/use-soft.htm>)

[xi] A few examples include the New York State Education and Research Network (NYSERNET), New England Academic and Research Network (NEARNET), the California Education and Research Foundation Network (CERFNET), Northwest Net (NWNET), Southern Universities Research and Academic Net (SURANET) and so on. UUNET was formed as a non-profit by a grant from the UNIX Users Group (USENIX).

[xii] UUNET called its Internet service ALTERNET. UUNET was acquired by Metropolitan Fiber Networks (MFS) in 1995 which was itself acquired by Worldcom in 1996. Worldcom later merged with MCI to form MCI WorldCom in 1998. In that same year, Worldcom also acquired the ANS backbone network from AOL, which had purchased it from the non-profit ANS earlier.

[xiii] PSINET was a for-profit spun out of the NYSERNET in 1990.

[xiv] CERFNET was started by General Atomics as one of the NSF-sponsored intermediate level networks. It was coincidental that the network was called "CERF"Net - originally they had planned to call themselves SURFNET, since General Atomics was located in San Diego, California, but this name was already taken by a Dutch Research organization called SURF, so the General Atomics founders settled for California Education and Research Foundation Network. Cerf participated in the launch of the network in July 1989 by breaking a fake bottle of champagne filled with glitter over a Cisco Systems router.

[xv] October 24, 1995, Resolution of the U.S. Federal Networking Council

RESOLUTION:

"The Federal Networking Council (FNC) agrees that the following language reflects our definition of the term "Internet".

"Internet" refers to the global information system that --

(i) is logically linked together by a globally unique address space based on the Internet Protocol (IP) or its subsequent extensions/follow-ons;

(ii) is able to support communications using the Transmission Control Protocol/Internet Protocol (TCP/IP) suite or its subsequent extensions/follow-ons, and/or other IP-compatible protocols; and

(iii) provides, uses or makes accessible, either publicly or privately, high level services layered on the communications and related infrastructure described herein."

[xvi] The Domain Name System was designed by Paul Mockapetris and initially documented in November 1983. Mockapetris, P., "Domain names - Concepts and Facilities", RFC 882, USC/Information Sciences Institute, November 1983 and Mockapetris, P., "Domain names - Implementation and Specification", RFC 883, USC/Information Sciences Institute, November 1983. (see also <http://soa.granitecanyon.com/faq.shtml>)

[xvii] The Handle System - see [www.handle.net](http://www.handle.net)

[xviii] See Leiner, et al, "A Brief History &", [www.isoc.org/internet/history/brief.html](http://www.isoc.org/internet/history/brief.html)

[xix] See [www.iana.org](http://www.iana.org) for more details. See also [www.icann.org](http://www.icann.org).

[xx] see [www.doi.org](http://www.doi.org)

[xxi] Version 5 of the Internet Protocol was an experiment which has since been terminated

[xxii] see A Framework for Distributed Digital Object Services, Robert E Kahn and Robert Wilensky at [www.cnri.reston.va.us/cstr/arch/k-w.html](http://www.cnri.reston.va.us/cstr/arch/k-w.html)

[xxiii] The interplanetary Internet effort is funded in part by DARPA and has

support from NASA. For more information, see [www.ipnsig.org](http://www.ipnsig.org)

[xxiv] See [www.iops.org](http://www.iops.org) for more information on this group dedicated to improving operational coordination among Internet Service Providers.