On the Evolution of Internet Technologies

VINTON G. CERF

Invited Paper

The Internet has been evolving from its origins in the early 1970s, based on work sponsored by the U.S. Defense Advanced Research Projects Agency. While the basic design was known in 1973 and first published in 1974 and the system essentially deployed in the academic and military communities on January 1, 1983, much has happened in the intervening 20 years. The first commercial Internet services emerged in 1989 after the interconnection of the Internet to commercial e-mail services. By 1993, commercial versions of the World Wide Web had appeared, and by 2003, voice over IP service was growing rapidly, after its first commercial introduction around 1995 (See Vocaltec: http://www.vocaltec.com/html/about/company.shtml).

The Internet of the future will be shaped by the tectonic forces of regulation, commercialization, technological change, and a wide range of policy concerns expressed at local, national, regional and international levels. In this paper, the effect of these forces is considered and an attempt made to project their effects into the future.

***Keywords*—*ARPA, BITNET, commercialization, Data networking, Defense Advanced Research Projects Agency (DARPA), domain names, European Academic Research Network (EARN), grid computing, history, information society, internationalized domain names, Internet, Internet Corporation for Assigned Names and Numbers (ICANN), Internet economics, Internet governance, Internet telephony, IP address space, IPv6, multiprotocol label switching, NetNorth, packet radio, packet satellite, packet switching, protocols, regulation, security, standards, telecommunications, X.25, virtual private network, Advanced Research Project Agency (ARPA), U.S. Department of Defense.*

I. Introduction

It is reasonable to ask what motivated the development of the Internet. In the early 1970s, one had little choice in networking. If one wanted to network computers from a particular vendor (e.g., IBM, Digital Equipment Corporation), one used networking technology that was proprietary to the vendor. IBM developed its Systems Network Architecture (SNA),[1] and Digital Equipment Corporation developed DECNET.[2] The U.S. Defense Department, having shown the utility of packet switching through its long-lived ARPANET project, pursued the idea of using computers in command and control. To avoid being constrained to a single vendor's equipment and networking technology, DARPA [1] set out in 1973 to develop a nonproprietary networking standard that would support computer-based command and control. It called the project Internetting [2] .

To understand the evolution of the Internet, one has to appreciate the basic architecture of the system and the ways in which it can and has evolved. As has been pointed out in many papers on the subject, the core of the Internet is the Internet Protocol (IP). Sometimes called the "thin waist" of the IP stack, the IP layer provides the basic glue that holds together the myriad networks of the Internet. It depends on a variety of other protocols to achieve this objective, not the least being several alternative routing protocols such as the *Border Gateway Protocol* (BGP and its variations), *Open Shortest Path First* (OSPF), *Internal System to Internet System* (IS-IS), and the relatively basic *Routing Information Protocol* (RIP).

In its earliest incarnations, the layering of the IPs had six layers beginning at the lowest level with the physical layer (e.g., Ethernet, SONET, T1, etc.), moving "up stack" through link (e.g., High Level Data Link Control (HDLC), Point-to-Point Protocol, etc.), Network (e.g., frame relay, asynchronous transfer mode (ATM), MPLS), Internet (e.g., IP, Internet Control Message Protocol), transport (e.g., Transmission Control Protocol, User Datagram Protocol), utility (e.g., File Transfer Protocol, Real Time Protocol, Simple Mail Transport Protocol, Post Office Protocol, Hypertext Transport Protocol), and application layers (e.g., e-mail clients, Web browsers, instant messaging clients). Subsequent formulations, some based on the seven-layer Open Systems Interconnection Model, incorporated the Internet layer in the network layer; some dropped the utility layer; some added session and/or presentation layers.

The original Internet design layered the IP on top of network layers implemented in ARPANET, the ARPA Packet Radio Network and the ARPA Atlantic Packet Satellite Network [1]. These systems were packet switching networks that encapsulated Internet packets as payload in their own packet

Manuscript received November 23, 2003; revised April 7, 2004.
The author is with MCI, Inc., Ashburn, VA 20147 USA.
Digital Object Identifier 10.1109/JPROC.2004.832974

[1][Online]. Available: http://www.yale.edu/pclt/COMM/SNA.HTM

[2][Online]. Available: http://en.wikipedia.org/wiki/DECnet

formats and routed the resulting ensemble to the next Internet *gateway* leading to the next network or to a destination *host*. As gateways became commercialized by pioneering companies such as Proteon[3] and Cisco Systems,[4] these devices came to be called *routers* because they operated on Internet packets and had to participate in protocols to propagate routing information at the IP level. As routers became the most common devices for implementing parts of the Internet, the notion of an underlying network diminished in favor of point-to-point or local area net links connecting the routers.

When frame relay and ATM networks emerged, they were referred to as *layer 2* networks, augmenting the notion of link layer with switching capability. Connections between the edge points of these networks were implemented as *virtual circuits*. More recently, a system called *multiprotocol label switching* (MPLS) has been developed which supports traffic-engineering requirements for IP networks and also, through stacking of labels, permits the creation of multiple virtual networks riding on top of the MPLS substrate.

The distinction between layers is somewhat blurred in these designs because the routing information needed to manage the MPLS substrate uses the same BGP as is used to support the IP layer. To make things even more complex, BGP itself relies on IP *and* the transport layer *Transmission Control Protocol* (TCP) for its implementation.

What is important to note about the architecture of the Internet is that the applications are far removed from the underlying transmission media. The IP separates the lower levels from the upper levels. In effect, the Internet layer is agnostic as to what means is used to transport Internet packets and is also agnostic as to what is carried in the payload of each Internet packet. A consequence of this agnosticism is that the Internet is capable of carrying virtually any digital content, including sound, voice, video, images, text, and so on. The implication is that services such as radio, television, and print publications can be transported through the Internet, assuming appropriate end-to-end transmission capacity. In many cases, these and related services do not require real-time support. For example, it is perfectly reasonable to transfer a high definition video as a kind of file transfer over the Internet, to be viewed at a later time. Of course, if the available data rate is low, "later" may be a long time. While real-time support for high quality video requires megabits per second, good quality audio can be carried in tens of kilobits per second. The latter puts audio services including voice communication well within the range of even dial-up access to the Internet.

The evolution of the Internet did not take place in a vacuum. Other parallel networking initiatives were undertaken in the academic and commercial sectors during the same period during which the main line Internet was developing. The commercial X.25 packet switching standard was developed in the mid-1970s by data networking groups in Canada, the United States, the United Kingdom, and France. The academic BITNET[5] emerged from academic networking of IBM mainframes and in part in response to the limited community of ARPANET players. BITNET used remote network job submission (NJE) as a means of moving data from one machine to another. Out of this effort emerged a robust e-mail list server system called LISTSERV and interactive chat capabilities, among many other things. BITNET users were linked to Internet users via e-mail, generally. BITNET had a European counterpart called the European Academic Research Network (EARN),[6] a Canadian counterpart, NetNorth, and a Japanese counterpart, AsiaNet, in Japan. BITNET was started in 1981, with its other counterparts emerging later. At its peak, BITNET was a globe-girdling system, but membership began to decline in 1993 as the commercial Internet and other academic networks overtook the BITNET technology.

Almost in parallel was the USENET,[7] based on the UNIX-to-UNIX Copy Program (UUCP), USENET emerged in 1981 as a kind of grassroots networking initiative based on the spread of the UNIX operating system. Out of this effort came the Net News service with its unique method of propagating information to many parties subscribing to different "feeds" associated with a huge variety of topics. Users could inject, comment upon and subscribe to specific topic areas, creating vast communities of users with mutual interest in a variety of topics. One of the founders of USENET went on to create a company, UUNET, that eventually became the operator of the largest Internet backbone. Today it is a part of MCI.

All of these networking systems were ultimately interconnected with the Internet. and as the Internet continued to spread, these services were layered on top of the increasingly ubiquitous TCP/IP protocols.

Evolution of telecommunications industry may be compared on several points to plate tectonics. While the concept of plate tectonics[8] is now a well-accepted fact, it is usually associated with very slow movement. The telecom industry is in the midst of a more cataclysmic change, but in some sense it is fair to compare the forces producing these changes to the kind of inevitability we associate with geological plate tectonics. These irresistible forces include technological change, regulation, commercial development, and policy evolution. While the combined effects of these forces cannot be predicted with precision, it seems fair to try to assess their near-term effects.

This paper aims to analyze the major aspects of the evolution of the Internet, from the lessons of the past to possible perspectives for the future. It is organized in three main sections. Section II deals with the numerous technical change

[3][Online]. Available: http://users.rcn.com/protn/services/

[4][Online]. Available: http://newsroom.cisco.com/dlls/company_overview.html

[5]Variously styled as "Because it's there" or "Because it's time" network. [Online]. Available: http://nethistory.dumbentia.com/; http://livinginternet.com/u/ui_bitnet.htm

[6][Online]. Available: http://www.cren.net/cren/cren-hist-fut.html; http://livinginternet.com/u/ui_bitnet.htm

[7][Online]. Available: http://www.cs.uu.nl/wais/html/na-dir/usenet/software/part1.html

[8][Online]. Available: http://pubs.usgs.gov/publications/text/dynamic.html

such as Virtual Private Networking, wireless communications, IP telephony or grid computing that pave the way of the Internet. Section III is dedicated to the commercial trends and forces of the Internet. Section IV discusses Internet governance, policy, and regulation.

II. Technological Change

One of the primary drivers of Internet evolution is technology. The Internet had its origins in the development of packet switching [1] in the 1960s and has continued to respond to new technological developments over the last 40 years. The introduction of commercial optical fiber communication in the 1980s in the form of synchronous optical networking (SONET/SDH) held the promise of vastly increased communication capacity for such networks. The introduction of frame relay, ATM switching, LANs, and, more recently, MPLS added to the mix of underlying transport and switching media over which the Internet can operate.

A. Virtual Private Networking

Virtual private networking has grown to become a major element in the use of networks in business, government, military and academic settings. The technologies for these networks have evolved in parallel with the Internet. ATM switching,[9] frame relay, and MPLS[10] are popular technologies for implementing virtual private networks. These technologies are also commonly used to carry encapsulated IP packets in networks that form part of the public Internet or are part of an enterprise virtual private network. MPLS, in particular, is becoming a popular technology for implementing public or private Internet services, providing traffic engineering and virtual private network separation through the judicious use of labels and distinct virtual routing tables. This technology is joined by earlier methods, such as encrypting IP in IP to form secure *tunnels* in the public Internet, to create enterprise virtual private networks. As the speed of IP and MPLS forwarding and switching increase, together with optical trunking capacity, these will become technologies of choice for broadband network users.

B. Wireless Communication

Wireless communication has transformed both voice telephony and data communications in the last decade. Mobile or cellular telephony has brought mobile communication to the developing and developed world in a dramatic way. Telegeography projects that there will be approximately 3 billion mobile phone subscribers by the end of 2005 and approximately 1.25 billion fixed line telephone subscribers.[11] The number of mobile phone subscribers in 1996 was only about 100 million. A significant evolution in wireless telephony is the addition of data in the form of digital transmission and short text messages in the Groupe Speciale Mobile (GSM) mobile telephone system.[12] GSM has been joined more recently by the code division multiple access (CDMA) technology for mobile use,[13] pioneered by Qualcomm Corporation,[14] among others. Systems using CDMA are sometimes referred to as third-generation (3G) systems. They provide data rates in the megabit- to multi-megabit-per-second range for digital communication.

In the data world, the most significant recent wireless change has been the widespread introduction of the *wireless LAN* based on the IEEE 802.11 standards.[15] Sometimes called *Wi-Fi* for *wireless fidelity*, 802.11 is one of 20 IEEE wireless or wired multiaccess network standards that have been developed in the last 35 years. The progenitor for these networks is the Alohanet developed in 1970 at the University of Hawaii, Manoa,[16] and which, itself, spawned the invention of Ethernet.[17] The data rates supported by these various developments reach into the scores of megabits per second. Newer developments include the so-called Wimax (IEEE 802.16)[18] and ultrawideband (UWB).[19] The pace of development of new wireless data technologies is significant and is joined by a similarly rapid development of wireless mobile communication.

Wireless mobile telephony is being extended by Internet enabling of wireless mobile devices. Wireless access to Internet through the IEEE 802 standards has spawned a generation of personal digital assistants (PDAs) that can communicate, as well as mobile laptop and notebook or notepad computers. Sensor systems are also being outfitted with wireless digital access to Internet service. A related wireless technology, called Bluetooth,[20] adds to the mix by eliminating the need for wires between devices in close proximity (e.g., keyboards, mice, telephone handsets, PDAs, and a variety of sensors). Bluetooth radios offer yet another means of transporting Internet traffic to wired access points.

As these systems evolve, it is reasonable to anticipate wider area coverage and higher capacities. For example, Vivato[21] has introduced a phased array technology for metropolitan area digital communication service. These systems employ switching and beam forming to provide service radii on the order of 4 km and data rates in the tens of megabits per second or more.

The wireless tidal wave is accompanied by new uses for Internet-enabled systems—for example, Internet-enabled automobiles.[22] It is already apparent that mobile or portable

[9][Online]. Available: http://www.atmforum.com

[10][Online]. Available: http://www.mplsforum.org/

[11][Online]. Available: http://www.telegeography.com/resources/statistics/telephony/intl_traffic_growth.html

[12][Online]. Available: http://ccnga.uwaterloo.ca/~jscouria/GSM/gsmreport.html

[13][Online]. Available: http://www.3gnewsroom.com/3g_news/nov_03/news_3948.shtml

[14][Online]. Available: http://www.qualcomm.com/about/index.html

[15][Online]. Available: http://grouper.ieee.org/groups/802/11/; http://grouper.ieee.org/groups/802/

[16][Online]. Available: http://en2.wikipedia.org/wiki/ALOHAnet

[17][Online]. Available: http://inventors.about.com/library/weekly/aa111 598.htm

[18][Online]. Available: http://www.wimaxforum.org/home

[19][Online]. Available: http://www.uwb.org/

[20][Online]. Available: http://www.bluetooth.com/

[21][Online]. Available: http://www.vivato.net/

[22][Online]. Available: http://www.vnunet.com/News/1 120 445

devices may have multiple avenues for connection to the Internet ranging from wired Ethernet to Bluetooth, 802.x and 3G alternatives. Some smart devices using software-defined radios[23] may scan the environment for the best choice of wireless connectivity, adapting as conditions change if the device is in fact mobile. A number of appliances are now Internet enabled—for example, the humble picture frame.[24]

One can anticipate the possibility that literally billions to tens or even hundreds of billions of devices will become Internet enabled as the integration of computing and communication continues.

C. IPv6

The current Internet uses version 4 of the IP (IPv4). While this has been sufficient, the 32-b address limits (4.3 billion addresses) have already spawned the use of private local addresses that have to be mapped into routable public IP addresses by means of network address translation devices (NATs). The Internet Engineering Task Force (IETF)[25] has developed a new standard, IPv6, which provides 128 b of address space (10^{38} addresses).[26] Progress on the implementation and deployment of IPv6 has been relatively slow but appears to be accelerating as more devices are developed with the capability to use this new IP. Some vendors, such as Sony, have already announced their intention to ship consumer devices in the near future with this capability.

The continuing integration of computing and communications is leading to multipurpose devices that serve as telephones, cameras, PDAs, electronic book readers, global positioning satellite receivers and so on. It is increasingly common for these devices to be Internet enabled.

D. Integration of Voice Telephony With Internet

One important development that is driving this integration is voice over IP (VOIP). A plethora of groups and technical developments surround this new capability[27] and have spawned significant production of devices to carry voice over the Internet or over IP networks and to interconnect these systems with the public switched telephone network. It was evident by 2003, from trade and news reports, that the telecom industry is adopting this technology in part out of demand from users and in part out of sheer competitive self-defense.

Among the new technologies contributing to the introduction of VOIP is the ENUM standard[28] that effectively allows international telephone numbers to be mapped into Internet domain names (actually, into so-called naming authority pointers). Internet enabled devices can look up a target telephone number in the Internet domain name system (DNS) and discover what Internet names and addresses are associated with it. This linkage makes it possible for a telephone call that originates in the public switched telephone network to be routed to an Internet termination or a call from an Internet originating device to be routed to a target in the Internet without passing through the public switched telephone network even though the target was referenced by telephone number.

Other technologies are critical to the implementation of VOIP, among which one has to include the Session Initiation Protocol (SIP)[29] and its various derivatives. SIP is used to "set up" and "tear down" voice calls that traverse the Internet. Another standard derived from the conventional telephony world is H323.[30] To some extent, these are competing protocols, but any provider of VOIP may have to implement both because existing, narrowband voice equipment is often equipped only with H323-compliant software. SIP is a very general protocol and can be used to implement general negotiations between communicating parties to establish the parameters that will guide and inform the protocols used to communicate. SIP could allow a supercomputer to negotiate with a PDA to determine the data rate and type of content that the PDA is capable of accepting or displaying, for example.

It is important to maintain a certain perspective about the use of the Internet to carry voice communications. This is simply *one* of many capabilities this versatile digital network can support. While we are quite deliberate about making phone calls today (dialing a number or even pushing to talk in a walkie–talkie environment), it is easy to see that voice communication may become simply a casual side effect of other modes of interaction.

Instant messaging has become an enormously popular tool for personal interaction on the Internet. It has gone from its origins as a consumer service to becoming an important part of the business world. One sees voice conferences augmented with instant messages among subsets of the conferees, for example. More important, the technology allows two people to begin a text conversation and migrate to voice mode or voice/video mode or even to shared *whitespace* mode in which a digital object, such as a PowerPoint file, might be displayed and edited by group collaboration. At no point in this sequence must a traditional phone call be made. It is this generalization of communications that makes the Internet such a powerful infrastructure and one which is demonstrably different from the communication systems that preceded it.

The topic of voice over the Internet or over IP will be revisited later in this paper in the context of the regulatory framework in which Internet services are considered.

E. A Multipurpose Internet

In general, it seems important to recognize that carrying voice over the Internet or over an IP backbone is simply one of myriad applications that the IP technology is capable of supporting. One can just as easily implement video conferencing and video and audio streaming (e.g., television and radio) over the Internet. Multicasting is one method for achieving these applications. There are two flavors of

[23][Online]. Available: http://www.sdrforum.org/

[24][Online]. Available: http://www.ceiva.com/

[25][Online]. Available: http://www.ietf.org

[26][Online]. Available: http://www.ipv6.org/

[27][Online]. Available: http://www.cis.ohio-state.edu/~jain/refs/ref_voip.htm

[28][Online]. Available: http://www.enum.org/information/resources.cfm

[29][Online]. Available: http://www.sipcenter.com/

[30][Online]. Available: http://www.packetizer.com/iptel/h323/

multicasting: single and multiple source. In the latter case, the Internet automatically calculates so-called spanning trees with built-in replication of packets to simulate what would normally have been a broadcast to all recipients. In the multiple-source case, every participant in the multicast group is capable of receiving *and sending* information to all other participants. In the case of single-source multicast, only one party can send, all others simply receive.

As broadband access to the Internet becomes more widely available, using technologies such as digital subscriber loops (DSL),[31] cable modems,[32] and digital satellite.[33] it becomes increasingly possible to support applications such as video conferencing, high-bandwidth group video gaming, and group collaboration. These capabilities contribute to the next major development for Internet, *grid computing*.[34]

F. Grid Computing

The concept here is deceptively simple: virtualize computing, storage, and communication resources in such a way that applications can simply acquire dynamic access to these resources to carry out a particular task and then return these resources to a pool for use in other applications. This is a kind of time sharing in three dimensions. To achieve this objective, the shared resources must not only be managed but also be made effectively fungible so that the repetition of a computation need not use precisely the same resources each time.

While it would be attractive to ignore all physical aspects of the distributed resources employed in a *grid*, reality dictates that speed-of-light propagation delays and other physical constraints, such as network communication capacity, may have to be taken into consideration in allocating resources to particular computations. For example, a prototypical example of a computing grid is found in the so-called search for extraterrestrial intelligence (SETI) application.[35] In this application, personal computers download an application that usually runs as a screen saver (that is, only when the machine is apparently idle). When this application is activated, it downloads a segment of a received radio signal from a predetermined source and runs a variety of analyses on the signal, looking for regularities that might indicate intelligent origin. Any "interesting" signals are reported back to the SETI central location. Because there is essentially no communication required among the millions of machines that might be running this application, this is an ideal application for gridlike treatment. A similar analysis shows that some kinds of cryptanalysis is also well suited to this style of computing. Simulations that do not require large amounts of intermediate data to be exchanged can also use this technique although the problem becomes harder the more intermediate data has to be exchanged among the computing elements in the course of the computation. Certain kinds of indexing operations might be candidates for grid treatment: each computing element is assigned a portion of the information space to be indexed and contributes its results to a location that combines the information into a single large database.

New protocols have been invented to support grid applications. At IBM these protocols are part of the WebSphere[36] system, and at Microsoft, they are part of the .NET[37] development. The protocols take advantage of protocols originally developed for the World Wide Web[38] and extended to become part of a suite of protocols known as *Web Services.*[39] These include *Web Services Description Language* (WSDL), *Extended Markup Language* (XML), *Simple Object Access Protocol* (SOAP), and *Universal Description and Discovery and Integration of Web Services* (UDDI),[40] among others.

Using these new protocols, it is possible to fabricate a virtually unlimited range of applications that will enable interactions among and between consumers, businesses, and governments. While it is still early in the evolution of Web Services, it seems clear that this technology will fuel a substantial opportunity for the creation of new products and services that operate over the Internet.

G. Internationalization of the DNS

The Internet DNS[41] was developed in the early to mid-1980s as a way of associating locations in the Internet with identifiers other than raw IP addresses. The system is highly distributed and resilient. Its hierarchical structure allows end users to manage the binding of a domain name with IP addresses by operating or outsourcing the operation of a so-called *name server*.

In the initial design of the system, domain names were limited to character strings using the Latin character set, including only letters A–Z, digits 0–9, and the dash ("—"). The hierarchical structure is denoted by separating symbols from different layers in the hierarchy with a period ("·"); for example, www.mci.com.

The rightmost string is known as a *top-level domain name* and the strings to the left are subdomains. There are 15 generic top level domains (.com, .net, .org, .gov, .int, .edu, .mil, .arpa, .coop, .aero, .pro, .name, .info, .museum, and .biz) and 243 country code top level domains (such as .us for United States and .za for South Africa).[42]

In a recent effort over the last 24 months, technologists in the IETF, among others, have worked to develop ways to encode character sets other than Latin characters into the DNS records. In particular, the so-called *unicode* character set[43] has been chosen as the primary reference for the scripts of many languages. These 16-b codes are further encoded

[31][Online]. Available: http://www.dslforum.org/

[32][Online]. Available: http://www.cablelabs.com/

[33][Online]. Available: http://www.direcpc.com/

[34][Online]. Available: http://www.gridforum.org/

[35][Online]. Available: http://setiathome.ssl.berkeley.edu/

[36]WebSphere is a trademark of the IBM Corporation.

[37].NET is a trademark of Microsoft Corporation.

[38][Online]. Available: http://www.w3.org/People/Berners-Lee/Weaving/Overview.html

[39][Online]. Available: http://www.w3.org/2002/ws/

[40][Online]. Available: http://www.uddi.org/

[41][Online]. Available: http://www.dns.net/dnsrd/, www.icann.org

[42][Online]. Available: http://www.iana.org/cctld/cctld-whois.htm

[43][Online]. Available: http://www.unicode.org/

through a careful process developed by the IETF and codified in the series of Internet standards documents known as requests for comments (RFCs) [3].

The effort to introduce these rich new non-Latin domain names into the DNS is still under way. There are many complex problems to solve in the process of accommodating this extension to the DNS, not the least of which is to determine so-called *restriction tables* for particular languages so that characters that would create duplicate names or would lead to user confusion can be minimized. In addition, the 16-b codes of the UNICODE system have to be encoded as valid 8-b character codes of the DNS that uses a restricted set of the American Standard Code for Information Interchange (ASCII) character set.

It is also important to recognize that the introduction of these non-Latin characters may have some negative side effects. If it is assumed that most Internet users are able to read and distinguish Latin characters from the restricted set A–Z, 0–9, and "—," that system of characters may be internationally recognizable, while characters from Mandarin, Arabic, Cyrillic, Georgian, or Hebrew, for example, may be less recognizable to everyone. A consequence of this is that the "internationalization" of the DNS may really be its "localization."

H. Security

A paper on Internet technology evolution would surely not be satisfactory if no mention were made of security. It is clear that the Internet is a popular system to attack. A variety of problems have surfaced in the last 15 years, including direct attacks against the Internet infrastructure by targeting hosts and routers or even elements of the DNS. Some of these attacks seek to exploit weaknesses in the operating systems of hosts, routers and infrastructure servers on the Internet. Some are variations on *denial-of-service* (DoS) attacks that direct large quantities of traffic to the target site. A variant of the latter is the so-called *distributed DoS attack (DDoS)* in which large numbers of hosts (typically PCs) are penetrated and equipped with *Trojan Horse* software that can be activated remotely to deliver large quantities of traffic to any targeted hosts. The first major attack against Internet hosts took place in 1988 and was dubbed a *worm* attack because the software that launched the attack replicated and transmitted itself around the Internet after compromising thousands of hosts running the UNIX operating system. In the wake of this attack, a Computer Emergency Response Team (CERT)[44] was created by the U.S. Defense Advanced Research Projects Agency (DARPA).

A variety of mechanisms have been developed to counter these attacks. Some involve the use of *firewalls* to observe traffic flowing in and out of an enterprise network and to restrict the use of certain protocols or to treat asymmetrically the establishment of protocol connections from *inside* the firewall and from the *outside.* Router vendors and security vendors have developed tools to detect and route attacking traffic to *blackholes* (so to speak) or to divert them for further analysis. *Intrusion detection systems* observe traffic patterns to try to detect various forms of attack.

Authentication is another important area in which significant development has occurred. The invention of *public-key cryptography* [4] was followed by the invention of a number of realizations of this system. Among the most well-known of these is the so-called *RSA algorithm*, invented by R. **R**ivest, A. **S**hamir, and L. **A**dleman. Out of these inventions came the concepts of *digital signatures* in addition to use of key pairs (*public key and private key)* to protect information while in transit or in storage.

Because public-key cryptography continues to be compute intensive, it is often the case that a conventional symmetric key (shared by both sender and receiver) are used to encrypt traffic, but the key itself is conveyed by one party to the other using public-key cryptography. Among the popular public symmetric methods for cryptography are the data encryption standard (DES),[45] developed by the U.S. National Institutes of Standards and Technology. A variant of this is the *triple-DES algorithm*[46] and the more recent *advanced encryption standard.*[47] Newer technologies are emerging, including so-called *elliptic codes*[48] that have been known for some time (30 years or more) but are just becoming available in unclassified applications.

All of these technologies foreshadow the possibility that strong authentication using strong cryptographic methods, possibly coupled with devices such as *smart cards*,[49] will become the technique of choice for identifying endpoints in networks, users of applications, participants in transactions, and the like.

Another area of intense security interest is labeled *authentication, access control, and accountability* (AAA). In this concept, parties are authenticated using passwords, or nonreusable passwords, or signed cryptographic certificates; once authenticated, databases containing their privileges (authorizations and access scope) can be consulted to determine what services and information they are entitled to access. Finally, to allow for the possibility of abuse of permissions, audit trails can be created for accountability. The AAA disciplines continue to develop and will become increasingly useful as standards for these mechanisms are adopted broadly. One of the "holy grails" in this area is the so-called single sign-on that would simplify user experiences in networked environments without compromising security.

III. Commercial Trends and Forces

Despite its long development history, commercial opportunities for the Internet did not emerge until the early 1980s. Several companies participated in early commercialization,

[44][Online]. Available: http://www.cert.org/

[45][Online]. Available: http://www.itl.nist.gov/fipspubs/fip46-2.htm

[46][Online]. Available: http://csrc.nist.gov/cryptval/des/tripledesval.html

[47][Online]. Available: http://csrc.nist.gov/CryptoToolkit/aes/

[48][Online]. Available: http://www.tcs.hut.fi/~helger/crypto/link/public/elliptic/

[49][Online]. Available: http://www.smartcardalliance.org/

including 3COM[50] and SUN Microsystems[51] in addition to Proteon and Cisco Systems. Proteon sold both token ring networks and Internet routers, and Cisco Systems focused on the latter. Commercial Internet services did not emerge until 1989 in the United States. It took some time after than before it emerged in commercial form in other parts of the world. In the United States, the ARPANET was retired in mid-1990, shortly after permission was given by the then U.S. Federal Networking Council[52] in 1988 to the Corporation for National Research Initiatives (CNRI)[53] to link the MCI Mail[54] commercial electronic mail system to the Internet. This interconnection took place in mid-1989 by way of a connection to the NSF Network (NSFNET).

In that same year, three commercial Internet Service Providers emerged: UUNET Technologies,[55] PSINet,[56] and CERFNET.[57] Since 1989, a significant number of ISPs have been founded, and, while many have failed, thousands remain, worldwide, providing Internet services to large and small customer bases.

A. Emerging Commercial Internet

The spread of the Internet, first to the academic communities starting in 1983 and then to commercial enterprise after 1989, fueled both the Internet equipment industry and the Internet service industry. Governments and military also contribute to the commercial growth of Internet products and services by using commercial off-the-shelf equipment, software, and services to satisfy many of their digital communication needs.

In 1994, the commercial World Wide Web emerged as Netscape Communications[58] began shipping its client and server software. Among its first customers was MCI Communications, which purchased licenses for clients and servers for Internet MCI commercial online services. The enormous success of Netscape's initial public offering in August 1995 fueled a boom in Internet-related businesses. Sometimes called the *dot-com boom*, venture capital rained in torrents on anyone with a business plan that mentioned the word *Internet.* Fortunes were made overnight and the explosion reached its peak in late 1999. By April 2000, the market had come to its senses and the great *dot-bomb* imploded. Fortunes were lost.

Throughout the period of Internet mania, however, there continued to be a consistent growth in the number of users of the Internet and the number of networks and machines connected to it. A prominent documenter of the real statistics behind Internet growth is A. Odlyzko, head of the Digital Technology Center, University of Minnesota, Minneapolis.[59] Odlyzko shows that while the Internet never sustained growth factors of 1000% per year, it has maintained a steady growth of 50%–80% per year since 1988. The rate of growth has diminished in the last year or two (to about 40–80%/year), but this is still a healthy pace and, as the world economy improves, this may even increase.

Setting aside the obvious commercial sale of devices and wholesale and retail Internet service, Internet-based businesses have continued to make their mark. Many fell by the wayside as the Internet dot-bomb struck, but those that survived appear to have sustainable business models. Many businesses invested heavily in enterprise Internet capabilities: internal Web sites, Web-based transaction processing, business-to-business and business-to-consumer Web-based services, and Web-based advertising. While the latter has not been the huge business it was once thought to be, companies such as Google[60] have succeeded in monetizing the search engine service, revolutionizing thinking about Web-based services. Others such as American OnLine (AOL)[61] and Yahoo![62] have survived. Many others have either been acquired or gone into chapter 11 or failed. In the period since April 2000, Internet oriented products and services have been through a widespread process of consolidation.

The relatively recent emergence of Web Services, mentioned in the section of this paper on technology, represents a new commercial focus on which a number of companies will depend for commercial growth.

B. Intellectual Property Protection

The general problem of *piracy*[63] has been a concern in the software industry for some years and has now reached book, movie and music publishers. The music industry has had a particularly visible involvement with the Internet—not entirely to its liking. As more music files were encoded in

[50]3COM was founded to make and sell Ethernets, and these proved enormously popular for academics interested in building LANs of powerful workstations.

[51]SUN Microsystems adopted TCP/IP and the rest of the IP suite as the basis for its networking, in contrast to IBM, which used its proprietary SNA; Xerox, which used Xerox Network System (XNS); Digital Equipment Corporation, which used DECNET; and Hewlett-Packard, which used its proprietary DS networking technology.

[52]Made up primarily of representatives from DARPA, the National Science Foundation (NSF), the National Aeronautics and Space Administration (NASA), the Department of Energy (DOE), the Department of Health and Human Services (HHS), and several smaller research-sponsoring organizations in the U.S. federal government.

[53][Online]. Available: http://www.cnri.reston.va.us

[54]MCI Mail is a trademark of MCI Corporation.

[55]UUNET Technologies was founded in 1987 as a nonprofit providing UNIX-to-UNIX-Copy Protocol service (UUCP) but was reformed as a for-profit in 1989. UUNET is now operating as part of the MCI Corporation's Internet service.

[56]PSINet was originally the nonprofit New York State Education and Research Network but spun out as a for-profit in 1989. PSINet later went bankrupt and its assets acquired by another company.

[57]CERFNET was founded in July 1989, by General Atomics. It was later acquired by TCG, which was itself acquired later by AT&T.

[58]Netscape Communications was founded in 1994 by J. Clark and M. Andreessen in a commercial reimplementation of the Mosaic World Wide Web browser developed at the National Center for Supercomputer Applications (NCSA) in 1992 by Andreessen and E. Bina.

[59][Online]. Available: http://www.dtc.umn.edu/~odlyzko/

[60][Online]. Available: http://www.google.com

[61]aol.com acquired Time-Warner during the period of AOL's most rapid growth; the company has renamed itself Time-Warner and AOL remains a division of the company.

[62][Online]. Available: http://www.yahoo.com

[63]Piracy is a term used by the intellectual property industry to characterize the theft of property, particularly digital property.

the MP3 format,[64] users began to accumulate and exchange these files. Eventually, a peer-to-peer method to exchange these files was developed by a company called Napster. The music industry reacted strongly to this development, ultimately suing Napster and its users for copyright infringement. This is just the tip of an iceberg in terms of intellectual property protection. Napster ultimately settled with the companies that sued it for infringement, but the same concerns arise with any digital content associated with intellectual property: digital movies, books, and software. Apple Computer introduced a commercial service it calls *iTunes*,[65] which permits users to download MP3-encoded music for a small fee. Since the service was inaugurated, an estimated 10 million songs have been downloaded. Apple released a hardware device it calls the *iPod*, and this device is used to download the music.

This is not the first time that technology has collided with an industry dependent on intellectual property protection. When Xerox first released the dry copier, book publishers were deeply concerned that they might sell one copy of a book and that the rest would be pirated. The same arguments were made about audio- and videocassette recorders, and now they are being made about digital copying via the Internet. In the case of the copiers, eventually the concept of *fair use* emerged that permitted individuals to duplicate parts or even all of a copyright work for personal uses (such as backing up a digital work on a removable storage medium such as a floppy disk or optical CD).

At issue is the quality of the digital recording (each copy is perfect or at least exactly as good as the original). These "perfect" recordings represent a major risk for the intellectual property owner, who may lose control over the work if it can be easily replicated. HDTV cable-casters and broadcasters have related concerns, especially with the arrival of digital home recorders such as the commercial TiVo system.[66] Disney recently engaged in trials of a service called MovieBeam,[67] in which protected files are transmitted over the air to a home receiver linked to a 160-GB disk drive. The drive holds about 100 movies and Disney plans to deliver up to ten movies a week to this storage system. The transmission is done over channels already accessible to Disney without interfering with the analog video signal that is *also* transmitted at the same time.

Just as these industries have ultimately adapted to new technologies that were thought to threaten them, it is reasonable to expect that the present tension over Internet propagation of intellectual property will be resolved. In the case of videocassettes, for example, the movie industry ultimately found that it could make three times as much money from the rental of prerecorded cassettes than it could make at the box office.

C. Broadband

Broadband access to the Internet continues to be a subject of great controversy. Everyone wants, it but not necessary under the same conditions. While the regulatory aspects of broadband will be explored in a later section, the commercial considerations still offer some important insights. There are at present only three relatively comparable options for broadband access today: cable modems, DSLs, and digital satellite. Other wireless alternatives such as multipoint microwave distribution service (MMDS) have not yet proven to be economically viable. Cable modem service outnumbers by about two to one the DSL service for consumers. Part of the reason is that customers are already likely committed to have cable service for television and see broadband Internet as an add-on. DSLs can be bundled with regular telephone service or Internet service, but do not have the benefit of entertainment television *pull*.

There continues to be controversy over the cost of broadband equipment and operation and under what circumstances it will be justified. The holy grail has tended to be optical fiber to the premises (business or residential). Business can often justify the cost of high bandwidth dedicated capacity because it is aggregating demand and it can presumably factor the cost into its business models. Residential broadband is harder to justify because it is an after-tax expense to the consumer. Despite these challenges, there appears to be continued demand for broadband access. The Federal Communications Commission (FCC) reported that 20 million homes had broadband service of some kind at the end of 2002. The U.S. population reached 291 500 000 in 2003,[68] and average family size declined to 2.6 persons per household in 2002,[69] suggesting an estimated 112 million households and, therefore, about 16% broadband penetration at the end of 2002. In 1999, only 2.8 million households had broadband, so the growth rate has been significant.

There is a great deal of speculation that the deployment of broadband will unlock an avalanche of services and economic investment. While there is reason to think that some applications require real-time or large file transfer (such as images or video or sound files), it should be noted that the available broadband residential technology tends to be asymmetric. That is, there is more capacity toward the user than available to the user to transmit. It would not be unreasonable to expect that personal information databases will continue to increase as storage media costs drop and grows in scale. The Disney MovieBeam service mentioned earlier involves a residential 160-GB disk drive, for example. This line of reasoning suggests that symmetry may be a necessary adjunct to broadband to trigger applications that involve serious data sharing among consumers. Of course, we have already seen

[64]MP3 means MPEG Audio layer 3. MPEG is the Motion Pictures Expert Group. It is an audio compression technology taken from the MPEG-1 and MPEG-2 specifications. MP3 compresses CD quality sound by a factor of 8 to 12, while maintaining almost the same high-fidelity sound quality. MP3 was developed by the German Frauenhofer Research Institute. Thomson Multimedia has patented MP3 in the United States and in Germany.

[65][Online]. Available: http://www.apple.com/itunes/

[66][Online]. Available: http://www.tivo.com/0.0.asp

[67][Online]. Available: http://www.internetnews.com/xSP/article.php/3 085 021

[68]U.S. Census Department Population Reference Bureau.

[69][Online]. Available: http://www.ameristat.org/Content/NavigationMenu/Ameristat/Topics1/MarriageandFamily/While_U_S__Households_Contract,_Homes_Expand.htm

some of the consumer behavior patterns with audio files even in the absence of symmetric broadband.

D. Other Consumer and Business Trends

While statistics were not available at the time of this writing, there is some indication that users are dropping second narrowband telephone lines in favor of wireless telephone service and/or broadband service (e.g., cable modems). There is also some indication that online bill presentation and payment is escalating. This should have an impact on the first class postal mail stream. Online shopping will increase the business of shipping companies such as FedEx, UPS, DHL, and the U.S. Postal Service, to name a few. Already companies like UPS are setting up warehouse sites to process orders directly from warehouse to delivery system as merchandisers outsource this service where it is cost effective.

Anecdotal reports also suggest that business-to-consumer transactions are on the increase as are business-to-business transactions, with the latter carrying higher value because of the size of individual transactions. Another clear trend is the adaptation of consumer relationship management to online modes. The tracking of shipments, flight status and trouble tickets, and order entry for reservations, books, flowers, and a host of other deliverables appear to be on the increase. The use of search engines to find specific products or services, including travel destinations, is augmenting if not replacing catalog sales.

On the whole, the rapidly increasing information content of the Internet is becoming an information infrastructure to which many turn first when in need of answers. As should be apparent from the footnotes in this paper, the Internet provided a substantial amount of reference information in the course of the writing.

IV. Internet Governance, Policy, and Regulation

This section is deserving of an entire book or perhaps a bookshelf of information. The time and space available for preparation limit what can be covered in this section, but it is vital to appreciate that policy issues surrounding the Internet are among the most difficult matters with which to deal. The issues are often global in scope, defying clear jurisdiction and potentially running afoul of mismatches in culture, laws, and social practices. At best, only some of the current issues will be addressed.

A. Internet Governance and ICANN

Internet governance is a slippery term that does not have a simple definition. It often means different things to different people and leads to heated debates as a consequence. One of the most visible policy bodies in the Internet space is the Internet Corporation for Assigned Names and Numbers (ICANN).[70] Created in October 1998, after substantial international debate that ultimately involved the U.S. White House, the European Commission, and many other interested parties, ICANN has a specific and relatively narrow mandate with regard to Internet policy. ICANN is responsible for overseeing the operation of the DNS and other unique identifiers needed for operation of the Internet, including the address space and a host of protocol parameters that must be maintained in well-defined tables for reference by software implementers and operators.

Since its creation, ICANN has struggled to limit the scope of its responsibilities, but as one of the few agencies charged with a policy development role for some aspect of Internet operation, aggrieved consumers, business people, and governments often turn to ICANN when they see no other place to go.

For purposes of this paper, *Internet governance* should *not* be defined to be constrained by the narrow limits of ICANN's mandate. Internet governance is a very broad topic and includes issues such as fraud, libel, slander, misrepresentation, taxation policy, intellectual property protection including trademark and copyright, consumer protection, operational policies associated with the stability and integrity of the DNS, privacy protection for personal information collected by domain name registrars and registries, and policies for the interception of e-mail, VOIP communication, instant messaging, and other communication modalities. Alert readers will be able to generate a much longer list and will appreciate that the scope of governance issues is very much broader than the specific responsibilities of ICANN.

Though the author is likely biased as a consequence of service as Chairman of the Board of ICANN, it seems important that ICANN not be forced to increase the scope of its responsibilities. It already has a significant mandate that is hard to fulfill. Rather, it will need to work with interested constituencies to find appropriate venues in which to cope with governance matters associated with the Internet. For example, there are disputes as to which party has the right to register a domain name in a given top-level domain. ICANN worked with the World Intellectual Property Organization (WIPO)[71] to develop a Uniform Domain Name Dispute Resolution Policy (UDRP)[72] that engages qualified parties to provide arbitration services to parties in dispute. WIPO provides one such service.[73] An important feature of the UDRP is that if either party is dissatisfied with the outcome of the arbitration, the terms and conditions of the arbitration specifically do not preclude litigation.

In matters concerning consumer complaints, ICANN turns to national bodies, where they exist, to address them. In the United States, this would be the Federal Trade Commission (FTC).[74] Criminal acts need the attention of law enforcement, which could range from local police or county sheriff to the national authorities. In the United States, the latter might involve the Federal Bureau of Investigation; the U.S. Marshals office; the U.S. Secret Service; the Bureau of Alcohol, Tobacco and Firearms; the Department of Justice; the Immigration and Naturalization Service; the Customs Service; the

[70] [Online]. Available: http://www.icann.org

[71] [Online]. Available: http://www.wipo.org/

[72] [Online]. Available: http://www.icann.org/udrp/udrp-policy-24oct99.htm

[73] [Online]. Available: http://arbiter.wipo.int/domains/index.html

[74] [Online]. Available: http://www.ftc.gov/

Drug Enforcement Administration; the Securities and Exchange Commission; the Treasury Department; and any of a number of other national-level organizations responsible for some form of law enforcement.

B. Regulation

Governance can also include regulation, and some aspects of the Internet may trigger concern from telecommunications regulatory agencies such as the U.S. FCC[75] or OFCOM in the United Kingdom.[76] One of the major difficulties associated with Internet governance is the extent to which it falls into previous regulatory regimes or requires a second look. It is tempting to argue that the use of the Internet for voice communication should follow the same regulatory regime that narrowband voice has used in the past, but this could prove to be a serious misstep. As has been suggested earlier, the Internet supports nearly all communication modalities—some concurrently. Rather than regulating it by summing up all the regulations for all the modalities it can carry, it seems important to reexamine the basis for regulation of older technology in the context of the Internet analogs.

In a refreshing departure from the norm, L. Solum has prepared a paper suggesting that regulation of the Internet, if any is needed, should follow the layered architecture of the Internet [5]. In the United States, the FCC is responsible for telecommunications regulation and the National Telecommunications and Information Agency[77] for some aspects of policy. The U.S. State Department represents the United States in treaty venues such as the International Telecommunications Union.[78] In other countries, the situation may be better or worse in terms of coordination among parts of the government.

With regard to the Solum thesis, regulation should be applied at the layer in the architecture closest to the layer at which the regulatory problem arises. Cross-layer regulation produces both under- and overregulation in the sense that some things are affected that should not be and others that should be are not. It is a kind of classical under- and overshoot. For example, a court in Pennsylvania ordered ISPs to block access to certain Web sites in Spain that were alleged to be carrying child pornography. Because the ISPs could not block specific pages, they were forced to block access to an entire IP address. There were tens of thousands of Web pages on the Spanish server that were *not* a problem, but all access to the server from the MCI UUNET backbone was blocked. Carrying out this blocking order at the ISP layer of protocol resulted in overblocking of Web content. Had it been possible to treat directly with the Web service provider (a Spanish company), the effect could have been made more precise.

Similarly, a blocking order was issued to the appropriate ISPs by a government that did not want to allow VOIP traffic to enter the country on the grounds that it took away international settlements revenue from the international carrier, which was owned partly by the government and partly by a private operator. Blocking at the ISP level proved maginally effective because the parties using VOIP could move to use other *logical ports* than those normally associated with VOIP applications. In this case the net was too narrow.

Solum's layered regulation and enforcement principles try to focus enforcement on the most appropriate layer in the Internet architecture.

In the United States, the FCC has sought to label VOIP an information service and not subject to conventional regulation. This position ignores a broadband reality: that the DSL or cable modem over with this traffic flows is layered and includes a transmission/transport layer that can be and should be subject of regulation if the parties offering the service have market power over what is sent over this bottleneck facility. Otherwise, the supplier of the broadband service might be able to bundle his offering in such a way that competition at higher layers of protocol could be effectively blocked.

Ideally one wants to stimulate competition among service providers, and the most assured method to achieve this would be to regulate the transport aspect of the broadband channel in such a way that all ISPs have equal access to the broadband infrastructure. This would open up cable modems and wi-fi to competition above the transport layer. This would simply re-create in the broadband world what we done in the narrowband switched voice service. Narrowband voice services include the possibility of dialing up any ISP; subscribers have unlimited choice as to ISP.

C. World Summit on the Information Society

The World Summit on the Information Society[79] is an International Telecommunication Union (ITU) and United Nations (UN) initiative to look broadly at information infrastructure and consider policy options that will stimulate its growth and accessibility around the world. This summit will take place in two stages. The first was held in mid-December 2003 in Geneva, and the second is scheduled for November 2004 in Tunisia. Among the many topics under discussion is the question of *Internet governance*. In the many debates leading up to the summit, it has been clear that the meaning and scope of Internet governance has been largely confounded in two ways. First, it is assumed that ICANN is doing Internet governance and that either: 1) it is too large a task for ICANN to accomplish and, therefore, that ICANN should relinquish all responsibility for all aspects of Internet operation or 2) ICANN's role is limited to its charter but that this is insufficient to cover all Internet governance issues and, therefore, the matter should be turned over to the UN or a treaty organization set up by the UN.

V. Conclusion

In summary, the Internet will continue to evolve as long as users have full access to the technical specifications of the Internet and have the opportunity to inject traffic into the Internet at any layer of the protocol, including the IP layer.

[75][Online]. Available: http://www.fcc.gov/

[76][Online]. Available: http://www.ofcom.org.uk/

[77][Online]. Available: http://www.ntia.doc.gov/

[78][Online]. Available: http://www.itu.int/home/

[79][Online]. Available: http://www.itu.int/wsis/

While this position creates some potential risk for the operators and other users of the Internet, it assures that there will be robust competition and opportunity for grass-roots-driven product and service development, including experiments conducted by graduate students and other researchers, where ever they may be. In ten years' time, the Internet may have 3 billion users and 30 billion devices online. Like every prediction, this one is likely also wrong, but it may actually fall short of the reality of Internet's seemingly inexorable penetration of the society of the 21st Century.

With regard to regulation, Solum's layered approach seems particularly well suited to the problem at hand. Regarding governance, ICANN should work with governments and nongovernmental organizations to identify ways in which the parties at interest may turn to other venues than the ICANN process for governance assistance. Nonetheless, ICANN does have some governance responsibilities and should actively seek to carry them out, while working with others to take up the tasks that fall outside ICANN's purview.

References

[1] B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. B. Postel, L. G. Roberts, and S. Wolff. A brief history of the Internet. Internet Soc., Reston, VA. [Online]. Available: http://www.isoc.org/internet/history/brief.shtml

[2] V. G. Cerf and R. E. Kahn, "A protocol for packet network interconnection," *IEEE Trans. Commun.*, vol. COM-22, pp. 627–641, May 1974.

[3] P. Fältström, P. Hoffman, and A. Costello. (2003) Request for comments 3490: Internationalizing domain names in applications (IDNA). [Online]. Available: ftp://ftp.rfc-editor.org/in-notes/rfc3490.txt

[4] W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Trans. Inform. Theory*, vol. IT-22, pp. 644–654, Nov. 1976.

[5] L. Solum. (2003, June) The layers principle: Internet architecture and the law. Loyola Law School, Los Angeles, CA. [Online]. Available: http://ssrn.com/abstract=416 263

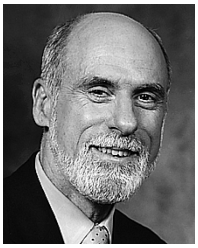

Vinton G. Cerf received the B.S. degree in mathematics from Stanford University, Stanford, CA, in 1965 and the M.A. and Ph.D. degrees in computer science from the University of California, Los Angeles (UCLA), in 1970 and 1972, respectively.

From 1982 to 1986, he led the engineering of MCI Mail, the first commercial e-mail service to be connected to the Internet. From 1986 to 1994, he was Vice-President of the Corporation for National Research Initiatives (CNRI). He is currently Senior Vice-President of Technology Strategy for MCI, Ashburn, VA. He is the codesigner of the TCP/IP protocols and the architecture of the Internet.

Dr. Cerf, with R. E. Kahn, received the U.S. National Medal of Technology for founding and developing the Internet in 1997. He is Honorary Chairman of the IPv6 Forum, served as Founding President of the Internet Society, and is Chairman of the Board of the Internet Corporation for Assigned Names and Numbers.